arXiv:1402.4157v2 [cs.AI] 12 May 2014

Conservative collision prediction and avoidance for stochastic trajectories in continuous time and space

Jan-P. Calliess, Michael Osborne, Stephen Roberts*
Dept. of Engineering Science, University of Oxford

May 13, 2014

Abstract

Existing work in multi-agent collision prediction and avoidance typically assumes discrete-time trajectories with Gaussian uncertainty or that are completely deterministic. We propose an approach that allows detection of collisions even between continuous, stochastic trajectories with the only restriction that means and covariances can be computed. To this end, we employ probabilistic bounds to derive criterion functions whose negative sign provably is indicative of probable collisions. For criterion functions that are Lipschitz, an algorithm is provided to rapidly find negative values or prove their absence. We propose an iterative policy-search approach that avoids prior discretisations and, upon termination, yields collision-free trajectories with adjustably high certainty. We test our method with both fixed-priority and auction-based protocols for coordinating the iterative planning process. Results are provided in collision-avoidance simulations of feedback controlled plants.

1 Introduction

Due to their practical importance, multi-agent collision avoidance and control have been extensively studied across different communities including AI, robotics and control. Considering continuous stochastic trajectories, reflecting each agent's uncertainty about its neighbours' time-indexed locations in an environment space, we exploit a distribution-independent bound on collision probabilities to develop a conservative collision-prediction module. It avoids

*The authours gratefully acknowledge funds via EPSRC EP/I011587.

temporal discretisation by stating collision-prediction as a one-dimensional optimization problem. If mean and standard deviation are computable Lipschitz functions of time, one can derive Lipschitz constants that allow us to guarantee collision prediction success with low computational effort. This is often the case, for instance, when dynamic knowledge of the involved trajectories is available (e.g. maximum velocities or even the stochastic differential equations).

To avoid collisions detected by the prediction module, we let an agent re-plan repeatedly until no more collisions occur with a definable probability. Here, re-planning refers to modifying a control signal (influencing the basin of attraction and equilibrium point of the agent's stochastic dynamics) so as to bound the collision probability while seeking low plan execution cost in expectation. To keep the exposition concrete, we focus our descriptions on an example scenario where the plans correspond to sequences of setpoints of a feedback controller regulating an agent's noisy state trajectory. However, one can apply our method in the context of more general policy search problems.

In order to foster low social cost across the entire agent collective, we compare two different coordination mechanisms. Firstly, we consider a simple fixed-priority scheme [11], and secondly, we modify an auction-based coordination protocol [7] to work in our continuous setting. In contrast to pre-existing work in auction-style multi-agent planning (e.g. [7,16]) and multi-agent collision avoidance (e.g. [1, 2, 15]), we avoid *a priori* discretizations of space and time. Instead, we recast the coordination problem as one of incremental open-loop policy search. That is, as a succession of continuous optimisation or root-finding problems that can be efficiently and reliably solved by modern optimisation and root-finding techniques (e.g. [13, 23]).

While our current experiments were conducted with linear stochastic differential equation (SDE) models with state-independent noise (yielding Gaussian processes), our method is also applicable to any situation where mean and covariances can be evaluated. This encompasses non-linear, non-Gaussian cases that may have state-dependent uncertainties (cf. [12]).

This preprint is an extended and improved version of a conference paper that appeared in *Proc. of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)* [6].

1.1 Related Work

Multi-agent trajectory planning and task allocation methods have been related to auction mechanisms by identifying locations in state space with atomic goods to be auctioned in a sequence of repeated coordination rounds (e.g. [7, 16, 26]). Unfortunately, even in finite domains the coordination is known to be intractable – for instance the sequential allocation problem is known to be NP-hard in the number of goods and agents [14, 22]. Furthermore, collision avoidance corresponds to non-convex interactions.

This renders the coordination problem inapplicable to standard optimization techniques that rely on convexity of the joint state space. In recent years, several works have investigated the use of mixed-integer programming techniques for

single- and multi-agent model-predictive control with collision avoidance both in deterministic and stochastic settings [7, 19]. To connect the problem to pre-existing mixed-integer optimization tools these works had to limit the models to dynamics governed by linear, time-discrete difference equations with state-independent state noise. The resulting plans were finite sequences of control inputs that could be chosen freely from a convex set. The controls gained from optimization are open-loop – to obtain closed-loop policies the optimization problems have to be successively re-solved on-line in a receding horizon fashion. However, computational effort may prohibit such an approach in multi-agent systems with rapidly evolving states.

Furthermore, prior time-discretisation comes with a natural trade-off. On the one hand, one would desire a high temporal resolution in order to limit the chance of missing a collision predictably occurring between consecutive time steps. On the other hand, communication restrictions, as well as poor scalability of mixed-integer programming techniques in the dimensionality of the input vectors, impose severe restrictions on this resolution. To address this trade-off, [10] proposed to interpolate between the optimized time steps in order to detect collisions occurring between the discrete time-steps. Whenever a collision was detected they proposed to augment the temporal resolution by the time-step of the detected collision thereby growing the state-vectors incrementally as needed. A detected conflict, at time t, is then resolved by solving a new mixed-integer linear programme over an augmented state space, now including the state at t. This approach can result in a succession of solution attempts of optimization problems of increasing complexity, but can nonetheless prove relatively computationally efficient. Unfortunately, their method is limited to linear, deterministic state-dynamics.

Another thread of works relies on dividing space into polytopes [1, 17], while still others [8,9,15,21] adopt a potential field. In not accommodating uncertainty and stochasticity, these approaches are forced to be overly conservative in order to prevent collisions in real systems.

In contrast to all these works, we will consider a different scenario. Our exposition focuses on the assumption that each agent is regulated by influencing its continuous stochastic dynamics. For instance, we might have a given feedback controller with which one can interact by providing a sequence of setpoints constituting the agent's plan. While this restricts the choice of control action, it also simplifies computation as the feedback law is fixed. The controller can generate a continuous, state-dependent control signal based on a discrete number of control decisions, embodied by the setpoints. Moreover, it renders our method applicable in settings where the agents' plants are controlled by standard off-the-shelf controllers (such as the omnipresent PID-controllers) rather than by more sophisticated customized ones. Instead of imposing discreteness, we make the often more realistic assumption that agents follow continuous time-state trajectories within a given continuous time interval. Unlike most work [1, 21, 25, 27] in this field, we allow for stochastic dynamics, where each agent cannot be certain about the location of its team-members. This is crucial for many real-world multi-agent systems. The uncertainties are modelled as state-noise which can re-

flect physical disturbances or merely model inaccuracies. While our exposition's focus is on stochastic differential equations, our approach is generally applicable in all contexts where the first two moments of the predicted trajectories can be evaluated for all time-steps. As noted above, this paper is an extended version of work that has been published in the proceedings of AAMAS'14 [6] and an earlier stage of this work was presented at an ICML [5] workshop.

2 Predictive Probabilistic Collision Detection with Criterion Functions

Task. Our aim is to design a collision-detection module that can decide whether a set of (predictive) stochastic trajectories is collision-free (in the sense defined below). The module we will derive is guaranteed to make this decision correctly, based on knowledge of the first and second order moments of the trajectories alone. In particular, no assumptions are made about the family of stochastic processes the trajectories belong to. As the required collision probabilities will generally have to be expressed as non-analytic integrals, we will content ourselves with a fast, *conservative* approach. That is, we are willing to tolerate a non-zero false-alarm-rate as long as decisions can be made rapidly and with zero false-negative rate. Of course, for certain distributions and plant shapes, one may derive closed-form solutions for the collision probability that may be less conservative and hence, lead to faster termination and shorter paths. In such cases, our derivations can serve as a template for the construction of criterion functions on the basis of the tighter probabilistic bounds.

Problem Formalization. Formally, a collision between two objects (or agents) $\mathfrak{a}, \mathfrak{r}$ at time $t \in I := [t_0, t_f] \subset \mathbb{R}$ can be described by the event

$\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t) = \{(x^{\mathfrak{a}}(t), x^{\mathfrak{r}}(t)) | \left\|x^{\mathfrak{a}}(t) - x^{\mathfrak{r}}(t)\right\|_2 \leq \frac{\Lambda^{\mathfrak{a}}+\Lambda^{\mathfrak{r}}}{2}\}$. Here, $\Lambda^{\mathfrak{a}}, \Lambda^{\mathfrak{r}}$ denote the objects' diameters, and $x^{\mathfrak{a}}, x^{\mathfrak{r}} : I \to \mathbb{R}^D$ are two (possibly uncertain) trajectories in a common, D-dimensional interaction space.

In a stochastic setting, we desire to bound the collision probability below a threshold $\delta \in (0, 1)$ at any given time in I. We loosely say that the trajectories are *collision-free* if $\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] < \delta, \forall t \in I$.

Approach. For conservative collision detection between two agents' stochastic trajectories $x^{\mathfrak{a}}, x^{\mathfrak{r}}$, we construct a *criterion function* $\gamma^{\mathfrak{a},\mathfrak{r}} : I \to \mathbb{R}$ (eq. as per Eq. 2 below). A conservative criterion function has the property $\gamma^{\mathfrak{a},\mathfrak{r}}(t) > 0 \Rightarrow \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] < \delta^{\mathfrak{a}}$. That is, a collision between the trajectories with probability above δ can be ruled-out if $\gamma^{\mathfrak{a},\mathfrak{r}}$ attains only positive values. If one could evaluate the function $t \mapsto \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)]$, an ideal criterion function would be

$$\gamma^{\mathfrak{a},\mathfrak{r}}_{\text{ideal}}(t) := \delta - \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)]. \tag{1}$$

It is ideal in the sense that $\gamma^{\mathfrak{a},\mathfrak{r}}_{\text{ideal}}(t) > 0 \Leftrightarrow \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] < \delta$. However, in most cases, evaluating the criterion function in closed form will not be feasible. Therefore, we adopt a conservative approach: That is, we determine a

criterion function $\gamma^{\mathfrak{a},\mathfrak{r}}(t)$ such that provably, we have $\gamma^{\mathfrak{a},\mathfrak{r}}(t) \leq \gamma^{\mathfrak{a},\mathfrak{r}}_{\text{ideal}}(t), \forall t$, including the possibility of false-alarms. That is, it is possible that for some times t, $\gamma^{\mathfrak{a},\mathfrak{r}}(t) \leq 0$, in spite of $\gamma^{\mathfrak{a},\mathfrak{r}}_{\text{ideal}}(t) > 0$.

Utilising the conservative criterion functions for collision-prediction, we assume a collision occurs unless $\min_{t\in I} \gamma^{\mathfrak{a},\mathfrak{r}}(t) > 0, \forall \mathfrak{r} \neq \mathfrak{a}$. If the trajectories' means and standard deviations are Lipschitz functions of time then one can often show that $\gamma^{\mathfrak{a},\mathfrak{r}}$ is Lipschitz as well. In such cases negative values of $\gamma^{\mathfrak{a},\mathfrak{r}}$ can be found or ruled out rapidly, as will be discussed in Sec. 2.1. In situations where a Lipschitz constant is unavailable or hard to determine, we can base our detection on the output of a global minimization method such as DIRECT [13].

2.1 Finding negative function values of Lipschitz functions

Let $t_0, t_f \in \mathbb{R}, t_0 \leq t_f, I := [t_0, t_f] \subset \mathbb{R}$. Assume we are given a Lipschitz continuous *target function* $f : I \to \mathbb{R}$ with Lipschitz constant $L \geq 0$. That is, $\forall S \subset I \, \exists L_S \leq L \, \forall x, x' \in S : |f(x) - f(x')| \leq L_S \, |x - x'|$. Let $t_0 < t_1 < t_2 < ... < t_N < t_f$ and define $G_N = (t_0, \ldots, t_{N+1})$ to be the *sample grid* of size $N + 2 \geq 2$ consisting of the inputs at which we choose to evaluate the target f.

Our goal is to prove or disprove the existence of a negative function value of target f.

2.1.1 A naive algorithm

As a first, naive method, Alg. 1 leverages Lipschitz continuity to answer the question of positivity correctly after a finite number of function evaluations.

The algorithm evaluates the function values on a finite grid assuming a uniform constant Lipschitz number L. The grid is iteratively refined until either a negative function value is found or, the Lipschitz continuity of function γ allows us to infer that no negative function values can exist. The latter is the case whenever $\min_{t\in G_N} \gamma(t) > L\,\Delta$ where $G_N = (t_0, ..., t_{N+1})$ is the grid of function input (time) samples, $\Delta = |t_{i+1} - t_i| \, (i = 0, ..., N - 1)$ and $L > 0$ a Lipschitz number of the function $\gamma : (t_0, t_f) \to \mathbb{R}$ which is to be evaluated.

The claim is established by the following Lemma:

Lemma 2.1. *Let $\gamma : [t_0, t_f] \subset \mathbb{R} \to \mathbb{R}$ be a Lipschitz function with Lipschitz number $L > 0$. Furthermore, let $G_N = (t_0, t_1, \ldots, t_{N+1})$ be an equidistant grid with $\Delta = |t_{i+1} - t_i| \, (i = 0, ..., N - 1)$.*

We have, $\gamma(t) > 0, \forall t \in (t_0, t_f)$ if $\forall t \in G_N : \gamma(t) > L\,\Delta$.

Proof. Since L is a Lipschitz constant of γ we have $|\gamma(t) - \gamma(t')| \leq L|t - t'|, \forall t, t' \in (t_0, t_f)$. Now, let $t^* \in (t_0, t_f)$ and $t_i, t_{i+1} \in G_N$ such that $t^* \in [t_i, t_{i+1}]$. Consistent with the premise of the implication we aim to show, we assume $\gamma(t_i), \gamma(t_{i+1}) > L\Delta$ and, without loss of generality, we assume $\gamma(t_i) \leq \gamma(t_{i+1})$. Let $\delta := |t_i - t^*|$. Since $t_i \leq t^* \leq t_{i+1}$ we have $0 \leq \Delta - \delta$. Finally, $0 < L\Delta < |\gamma(t_i)|$ implies $\gamma(t^*) \geq \gamma(t_i) - |\gamma(t_i) - \gamma(t^*)| \geq \gamma(t_i) - L|t_i - t^*| > L\Delta - L\delta = L(\Delta - \delta) \geq 0$. □

input : Domain boundaries $t_0, t_f \in \mathbb{R}$, function $\gamma : (t_0, t_f) \to \mathbb{R}$, Lipschitz constant $\mathsf{L} > 0$.

output: Flag flag indicating presence of a non-positive function value ($\mathsf{flag} = 1$ indicates existence of a non-positive function value; $\mathsf{flag} = 0$ indicates it has been ruled out that a negative function value can exist). Variable $\mathsf{criticalTime}$ contains the time of a non-positive function value if such exists ($\mathsf{criticalTime} = t_0 - 1$, iff $\gamma((t_0, t_f)) \subset \mathbb{R}_+$).

$\mathsf{flag} \leftarrow -1$; $\mathsf{criticalTime} \leftarrow t_0 - 1$;
$\mathsf{TimeGrid} \leftarrow \{t_0, t_f\}$;
$r \leftarrow -1$;
repeat
  $r \leftarrow r + 1$; $\Delta \leftarrow \frac{t_f - t_0}{2^r}$; $N \leftarrow (t_f - t_0)/\Delta$;
  $\mathsf{TimeGrid} \leftarrow \cup_{i=0}^{N} \{t_0 + i\Delta\}$;
  $\mathsf{minVal} \leftarrow \min_{t \in \mathsf{TimeGrid}} \gamma(t)$;
  if $\mathsf{minVal} \leq 0$ **then**
    $\mathsf{flag} \leftarrow 1$; $\mathsf{criticalTime} \leftarrow \arg\min_{t \in \mathsf{TimeGrid}} \gamma(t)$;
  else if $\mathsf{minVal} > \mathsf{L}\,\Delta$ **then**
    $\mathsf{flag} \leftarrow 0$;
until $\mathsf{flag} = 1$ OR $\mathsf{flag} = 0$;

Algorithm 1: Naive algorithm deciding whether a Lipschitz continuous function γ has a non-positive value on a compact domain. Note, if minVal $>$ L Δ the function is guaranteed to map into the positive reals exclusively.

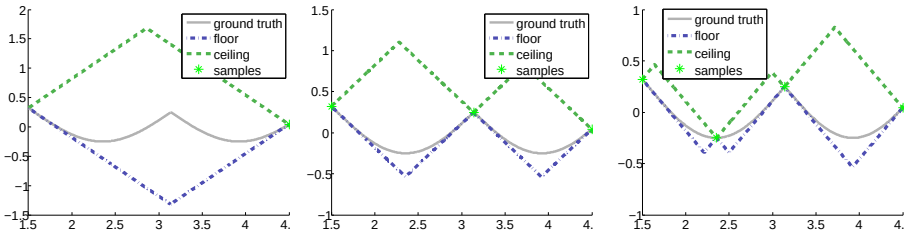

Figure 1: Proving the existence of a negative value of function $x \mapsto |\sin(x)|\cos(x) + \frac{1}{4}$. Left: Initial condition. Centre: First refinement. Right: The second refinement has revealed the existence of a negative value.

Appart from a termination criterion, the lemma establishes that larger Lipschitz numbers will generally cause longer run-times of the algorithm as finer resolutions Δt will be required to ensure non-negativity of the function under investigation.

2.1.2 An improved adaptive algorithm

Next, we will present an improved version of the algorithm provided above. We can define two functions, *ceiling* $\mathfrak{u}_N$ and *floor* $\mathfrak{l}_N$, such that (i) they bound the target $\forall t \in I : \mathfrak{l}_N(t) \leq \gamma(t) \leq \mathfrak{u}_N(t)$, and (ii) the bounds get tighter for denser grids. In particular, one can show that $\mathfrak{l}_N, \mathfrak{u}_N \overset{N\to\infty}{\longrightarrow} f$ uniformly if G_N converges to a dense subset of $[a, b]$. Define $\xi^{\mathfrak{l}}_N := \arg\min_{x\in I} \mathfrak{l}_N(x)$. It has been shown that $\xi^{\mathfrak{l}}_N = \min_{i=1}^{N-1} \frac{t_{i+1}+t_i}{2} - \frac{\gamma(t_{i+1})-\gamma(t_i)}{2L}$ and $\mathfrak{l}_N(\xi^{\mathfrak{l}}_N) = \min_i \frac{\gamma(t_{i+1})+\gamma(t_i)}{2} - L\frac{t_{i+1}-t_i}{2}$ (see [13, 23]). It is trivial to refine this to take localised Lipschitz constants into account: $\xi^{\mathfrak{l}}_N = \min_i \frac{\gamma(t_{i+1})+\gamma(t_i)}{2} - L_{J_i}\frac{t_{i+1}-t_i}{2}$ where L_{J_i} is a Lipschitz number valid on interval $J_i = (t_i, t_{i+1})$.

This suggests the following algorithm: *We refine the grid G_N to grid G_{N+1}, by including $\xi^{\mathfrak{l}}_N, f(\xi^{\mathfrak{l}}_N)$ as a new sample. This process is repeated until either of the following stopping conditions are met: (i) a negative function value of γ is discovered ($f(\xi^{\mathfrak{l}}_N) < 0$), or (ii) $\mathfrak{l}_N(\xi^{\mathfrak{l}}_N) \geq 0$ (in which case we are guaranteed that no negative function values can exist).*

For pseudo-code refer to Alg. 2.

An example run is depicted in Fig. 1. Note, without our stopping criteria, our algorithm degenerates to Shubert's minimization method [23]. The stopping criteria are important to save computation, especially in the absence of negative function values.

2.2 Deriving collision criterion functions

This subsection is dedicated to the derivation of a (Lipschitz) criterion function. In lieu to the approach of [7, 20], the idea is to define hyper-cuboids $H^{\mathfrak{a}}, H^{\mathfrak{r}}$ sufficently large to contain a large enough proportion of each agent's probability mass to ensure that no collision occurs (with sufficient confidence) as long as the

input : Domain boundaries $t_0, t_f \in \mathbb{R}$, function $\gamma : (t_0, t_f) \to \mathbb{R}$, Lipschitz constant $\mathsf{L} > 0$.

output: Flag flag indicating presence of a non-positive function value ($\mathsf{flag} = 1$ indicates existence of a non-positive function value; $\mathsf{flag} = 0$ indicates it has been ruled out that a negative function value can exist). Variable $\mathsf{criticalTime}$ contains the time of a non-positive function value if such exists ($\mathsf{criticalTime} = t_0 - 1$, iff $\gamma((t_0, t_f)) \subset \mathbb{R}_+$).

$\mathsf{flag} \leftarrow -1$; $\mathsf{criticalTime} \leftarrow t_0 - 1$;
$G_N \leftarrow \{t_0, t_f\}$;
$N = 0$;
repeat
  $\xi^{\mathfrak{l}} \leftarrow \min_{i=1}^{N} \frac{t_{i+1}+t_i}{2} - \frac{\gamma(t_{i+1})-\gamma(t_i)}{2L}$;
  $\mathfrak{l}_N(\xi_N^{\mathfrak{l}}) \leftarrow \min_{i=1}^{N} \frac{\gamma(t_{i+1})+\gamma(t_i)}{2} - L\frac{t_{i+1}-t_i}{2}$;
  $\mathsf{minVal} \leftarrow \gamma(\xi^{\mathfrak{l}})$;
  if $\mathsf{minVal} \leq 0$ **then**
    $\mathsf{flag} \leftarrow 1$; $\mathsf{criticalTime} \leftarrow \xi^{\mathfrak{l}}$;
  else if $\mathfrak{l}_N(\xi_N^{\mathfrak{l}}) > 0$ **then**
    $\mathsf{flag} \leftarrow 0$;
  else
    $N \leftarrow N + 1$; $G_N \leftarrow G_N \cup \{\xi^{\mathfrak{l}}\}$;
  end
until $\mathsf{flag} = 1$ OR $\mathsf{flag} = 0$;

Algorithm 2: Adaptive algorithm based on Shubert's method to prove whether a Lipschitz continuous function γ has a non-positive value on a compact domain. Note, if $\mathfrak{l}_N(\xi_N^{\mathfrak{l}}) > 0$ the function is guaranteed to map into the positive reals exclusively.

cuboids do not overlap. We then define the criterion function so as to negative values whenever the hyper-cuboids do overlap.

For ease of notation, we omit the time index t. For instance, in this subsection, $x^{\mathfrak{a}}$ now denotes random variable $x^{\mathfrak{a}}(t)$ rather than the stochastic trajectory.

The next thing we will do is to derive sufficient conditions for absence of collisions, i.e. for $\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}] < \delta$.

To this end, we make an intermediate step: For each agent $\mathfrak{q} \in \{\mathfrak{a}, \mathfrak{r}\}$ we define an open hyper-cuboid $H^{\mathfrak{q}}$ centred around mean $\mu^{\mathfrak{q}} = \langle x^{\mathfrak{q}}(t) \rangle$. As a D-dimensional hyper-cuboid, $H^{\mathfrak{q}}$ is completely determined by its centre point $\mu^{\mathfrak{q}}$ and its edge lengths $l_1^{\mathfrak{q}}, ..., l_D^{\mathfrak{q}}$. Let $O^{\mathfrak{q}}$ denote the event that $x^{\mathfrak{q}} \notin H^{\mathfrak{q}}$ and $P^{\mathfrak{q}} := \Pr[O^{\mathfrak{q}}]$. We derive a simple disjunctive constraint on the component distances of the means under which we can guarantee that the collision probability is not greater than the probability of at least one object being outside its hyper-cuboid. This is the case if the hypercuboids do not overlap. That is, their max-norm distance is at least $\Lambda^{\mathfrak{a},\mathfrak{r}} := \frac{\Lambda^{\mathfrak{a}} + \Lambda^{\mathfrak{r}}}{2}$.

Before engaging in a formal discussion we need to establish a preparatory fact:

Lemma 2.2. *Let $\mu_j^{\mathfrak{q}}$ denote the jth component of object $\mathfrak{q}$'s mean and $r_j^{\mathfrak{q}} = \frac{1}{2} l_j^{\mathfrak{q}}$. Furthermore, let $\mathfrak{F}^{\mathfrak{a},\mathfrak{r}} := \overline{\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}}$ be the event that no collision occurs and $\mathfrak{B}^{\mathfrak{a},\mathfrak{r}} := H^{\mathfrak{a}} \times H^{\mathfrak{r}}$ the event that $x^{\mathfrak{a}} \in H^{\mathfrak{a}}$ and $x^{\mathfrak{r}} \in H^{\mathfrak{r}}$. Assume the component-wise distance between the hyper-cuboids $H^{\mathfrak{a}}, H^{\mathfrak{r}}$ is at least $\Lambda^{\mathfrak{a},\mathfrak{r}}$, which is expressed by the following disjunctive constraint:*

$$\exists j \in \{1, ..., D\} : \left| \mu_j^{\mathfrak{a}} - \mu_j^{\mathfrak{r}} \right| > \Lambda^{\mathfrak{a},\mathfrak{r}} + r_j^{\mathfrak{a}} + r_j^{\mathfrak{r}}.$$

Then, we have : $\mathfrak{B}^{\mathfrak{a},\mathfrak{r}} \subset \mathfrak{F}^{\mathfrak{a},\mathfrak{r}}$.

Proof. Since $\|x\|_\infty \leq \|x\|_2, \forall x$ we have

$\mathfrak{F}_\infty := \{(x^{\mathfrak{a}}, x^{\mathfrak{r}}) | \, \|x^{\mathfrak{a}} - x^{\mathfrak{r}}\|_\infty > \Lambda^{\mathfrak{a},\mathfrak{r}}\} \subset \{(x^{\mathfrak{a}}, x^{\mathfrak{r}}) | \, \|x^{\mathfrak{a}} - x^{\mathfrak{r}}\|_2 > \Lambda^{\mathfrak{a},\mathfrak{r}}\} = \mathfrak{F}^{\mathfrak{a},\mathfrak{r}}$. It remains to be shown that $\mathfrak{B}^{\mathfrak{a},\mathfrak{r}} \subset \mathfrak{F}_\infty$: Let $(x^{\mathfrak{a}}, x^{\mathfrak{r}}) \in \mathfrak{B}^{\mathfrak{a},\mathfrak{r}} = H^{\mathfrak{a}} \times H^{\mathfrak{a}}$. Thus, $\forall j \in \{1, ..., D\}, \mathfrak{q} \in \{\mathfrak{a}, \mathfrak{r}\} : \left| x_j^{\mathfrak{q}} - \mu_j^{\mathfrak{q}} \right| \leq r_j^{\mathfrak{q}}$. For contradiction, assume $(x^{\mathfrak{a}}, x^{\mathfrak{r}}) \notin \mathfrak{F}_\infty$. Then, $|x_i^{\mathfrak{a}} - x_i^{\mathfrak{r}}| \leq \Lambda^{\mathfrak{a},\mathfrak{r}}$ for all $i \in \{1, ..., D\}$.

Hence, $|\mu_i^{\mathfrak{a}} - \mu_i^{\mathfrak{r}}| = |\mu_i^{\mathfrak{a}} - x_i^{\mathfrak{a}} + x_i^{\mathfrak{a}} - x_i^{\mathfrak{r}} + x_i^{\mathfrak{r}} - \mu_i^{\mathfrak{r}}| \leq |\mu_i^{\mathfrak{a}} - x_i^{\mathfrak{a}}| + |x_i^{\mathfrak{a}} - x_i^{\mathfrak{r}}| + |x_i^{\mathfrak{r}} - \mu_i^{\mathfrak{r}}| \leq r_i^{\mathfrak{a}} + \Lambda^{\mathfrak{a},\mathfrak{r}} + r_i^{\mathfrak{r}}, \forall i \in \{1, ..., D\}$ which contradicts our disjunctive constraint in the premise of the lemma. q.e.d.

□

Theorem 2.3. *Let $\mu_j^{\mathfrak{q}}$ denote the jth component of object $\mathfrak{q}$'s mean and $r_j^{\mathfrak{q}} = \frac{1}{2} l_j^{\mathfrak{q}}$. Assume, $x^{\mathfrak{a}}, x^{\mathfrak{r}}$ are random variables with means $\mu^{\mathfrak{a}} = \langle x^{\mathfrak{a}} \rangle, \mu^{\mathfrak{r}} = \langle x^{\mathfrak{r}} \rangle$, respectively. The max-norm distance between hypercuboids $H^{\mathfrak{a}}, H^{\mathfrak{r}}$ is at least $\Lambda^{\mathfrak{a},\mathfrak{r}} > 0$ (i.e. the hypercuboids do not overlap), which is expressed by the following disjunctive constraint:*

$$\exists j \in \{1, ..., D\} : \left| \mu_j^{\mathfrak{a}} - \mu_j^{\mathfrak{r}} \right| > \Lambda^{\mathfrak{a},\mathfrak{r}} + r_j^{\mathfrak{a}} + r_j^{\mathfrak{r}}.$$

Then, we have :

$$\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}] \leq P^{\mathfrak{a}} + P^{\mathfrak{r}} - P^{\mathfrak{r}} P^{\mathfrak{a}} \leq P^{\mathfrak{a}} + P^{\mathfrak{r}}$$

where $P^{\mathfrak{q}} = \Pr[x^{\mathfrak{q}} \notin H^{\mathfrak{q}}], (\mathfrak{q} \in \{\mathfrak{a}, \mathfrak{r}\})$.

Proof. As in Lem. 2.2, let $\mathfrak{F}^{\mathfrak{a},\mathfrak{r}} := \overline{\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}}$ be the event that no collision occurs and let $\mathfrak{B}^{\mathfrak{a},\mathfrak{r}} := H^{\mathfrak{a}} \times H^{\mathfrak{r}}$. We have $\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}] \leq 1 - \Pr[\overline{\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}}] = 1 - \Pr[\mathfrak{F}^{\mathfrak{a},\mathfrak{r}}]$. By Lem. 2.2 we have $\mathfrak{B}^{\mathfrak{a},\mathfrak{r}} \subset \mathfrak{F}^{\mathfrak{a},\mathfrak{r}}$ and thus, $1 - \Pr[\mathfrak{F}^{\mathfrak{a},\mathfrak{r}}] \leq 1 - \Pr[\mathfrak{B}^{\mathfrak{a},\mathfrak{r}}] = \Pr[\overline{\mathfrak{B}^{\mathfrak{a},\mathfrak{r}}}]$. Now, $\Pr[\overline{\mathfrak{B}^{\mathfrak{a},\mathfrak{r}}}] = \Pr[x^{\mathfrak{a}} \notin H^{\mathfrak{a}} \vee x^{\mathfrak{r}} \notin H^{\mathfrak{r}}] = P^{\mathfrak{a}} + P^{\mathfrak{r}} - P^{\mathfrak{a}}\,P^{\mathfrak{r}} \leq P^{\mathfrak{a}} + P^{\mathfrak{r}}$. q.e.d. □

One way to define a criterion function is as follows:

$$\gamma^{\mathfrak{a},\mathfrak{r}}(t; \varrho(t)) := \max_{i=1,\ldots,D} \{|\mu_i^{\mathfrak{a}}(t) - \mu_i^{\mathfrak{r}}(t)| - \Lambda^{\mathfrak{a},\mathfrak{r}} - r_i^{\mathfrak{a}}(t) - r_i^{\mathfrak{r}}(t)\} \tag{2}$$

where $\varrho = (r_1^{\mathfrak{a}}, \ldots, r_D^{\mathfrak{a}}, r_1^{\mathfrak{r}}, \ldots, r_D^{\mathfrak{r}})$ is the parameter vector of radii. (For notational convenience, we will often omit explicit mention of parameter ϱ in the function argument.)

For more than two agents, agent $\mathfrak{a}'s$ overall criterion function is $\Gamma^{\mathfrak{a}}(t) := \min_{\mathfrak{r} \in \mathfrak{A} \setminus \{\mathfrak{a}\}} \gamma^{\mathfrak{a},\mathfrak{r}}(t)$.

Thm. 2.3 tells us that the collision probability is bounded from above by the desired threshold δ if $\gamma^{\mathfrak{a},\mathfrak{r}}(t) > 0$, provided we chose the radii $r_j^{\mathfrak{a}}, r_j^{\mathfrak{r}}$ $(j = 1, ..., D)$ such that $P^{\mathfrak{a}}, P^{\mathfrak{r}} \leq \frac{\delta}{2}$.

Let $\mathfrak{q} \in \{\mathfrak{a}, \mathfrak{r}\}$. Probability theory provides several distribution-independent bounds relating the radii of a (possibly partly unbounded) hypercuboid to the probability of not falling into it. That is, these are bounds of the form

$$P^{\mathfrak{q}} \leq \beta(r_1^{\mathfrak{q}}, ..., r_D^{\mathfrak{q}}; \Theta)$$

where β is a continuous function that decreases monotonically with increasing radii and Θ represents additional information. In the case of Chebyshev-type bounds information about the first two moments are folded in, i.e. $\Theta = (\mu^{\mathfrak{q}}, C^{\mathfrak{q}})$ where $C^{\mathfrak{q}}(t) \in \mathbb{R}^{D \times D}$ is the variance (-covariance) matrix. We then solve for radii that fulfil the inequality $\frac{\delta}{2} \geq \beta(r_1^{\mathfrak{q}}, ..., r_D^{\mathfrak{q}}; \Theta)$ while simultaneously ensuring collision avoidance with the desired probability.

Inspecting Eq. 2, it becomes clear that, in order to maximally diminish conservatism of the criterion function, it would be ideal to choose the radii in ϱ such that $\varrho = \text{argmax}_{\varrho} \gamma^{\mathfrak{a},\mathfrak{r}}(t; \varrho) = \text{argmax}_{r_1^{\mathfrak{a}}, \ldots, r_D^{\mathfrak{a}} r_1^{\mathfrak{r}}, \ldots, r_D^{\mathfrak{r}}} \max_{i=1,\ldots,D} \{|\mu_i^{\mathfrak{a}} - \mu_i^{\mathfrak{r}}| - \Lambda^{\mathfrak{a},\mathfrak{r}} - r_i^{\mathfrak{a}} - r_i^{\mathfrak{r}}\}$ subject to the constraints $\frac{\delta}{2} \geq \beta(r_1^{\mathfrak{q}}, ..., r_D^{\mathfrak{q}}; \Theta), (\mathfrak{q} \in \{\mathfrak{a}, \mathfrak{r}\})$. Solving this constrained optimisation problem can often be done in closed form.

In the context where β is derived from a Chebyshev-type bound, we propose to set as many radii as large as possible (in order to decrease (β to satisfy the constraints) while setting the radii $r_i^{\mathfrak{a}}, r_i^{\mathfrak{r}}$ as small as possible without violating the constraint (where i is some dimension). That is, we define the radii as follows: Set $r_j^{\mathfrak{q}} := \infty, \forall j \neq i$. The remaining unknown variable, $r_i^{\mathfrak{q}}$, then is defined as the solution to the equation $\frac{\delta}{2} = \beta(r_1^{\mathfrak{q}}, ..., r_D^{\mathfrak{q}}; \Theta)$. The resulting criterion function, denoted by $\gamma_i^{\mathfrak{a},\mathfrak{r}}$, we obtain with this procedure of course depends on the arbitrary choice of dimension i. Therefore, we obtain a less conservative criterion function by repeating this process for each dimension i and then constructing a new criterion function as the point-wise maximum: $\gamma^{\mathfrak{a},\mathfrak{r}}(t) := \max_i \gamma_i^{\mathfrak{a},\mathfrak{r}}(t)$.

A concrete example of this procedure is provided below.

2.2.1 Example constructions of distribution-independent criterion functions

We can use the above derivation as a template for generating criterion functions.

Consider the following concrete example. Combining union bound and the standard (one-dim.) Chebyshev bound yields $P^{\mathfrak{q}} = \Pr[x^{\mathfrak{q}} \notin H^{\mathfrak{q}}] \leq \sum_{j=1}^{D} \frac{C^{\mathfrak{q}}_{jj}}{r^{\mathfrak{q}}_j r^{\mathfrak{q}}_j} =: \beta(r^{\mathfrak{q}}_1, \ldots, r^{\mathfrak{q}}_D; C^{\mathfrak{q}})$. Setting every radius, except $r^{\mathfrak{q}}_i$, to infinitely large values and β equal to $\frac{\delta}{2}$ yields $\frac{\delta}{2} = \frac{C^{\mathfrak{q}}_{ii}}{r^{\mathfrak{q}}_i r^{\mathfrak{q}}_i}$, i.e. $r^{\mathfrak{q}}_i = \sqrt{\frac{2C^{\mathfrak{q}}_{ii}}{\delta}}$. (Note, this a correction of the radius provided in the conference version of this paper.) Finally, inserting these radii (for $\mathfrak{q} = \mathfrak{a}, \mathfrak{r}$) into Eq. 2 yields our first collision criterion function: $\gamma^{\mathfrak{a},\mathfrak{r}}(t) := |\mu^{\mathfrak{a}}_i(t) - \mu^{\mathfrak{r}}_i(t)| - \Lambda^{\mathfrak{a},\mathfrak{r}} - \sqrt{\frac{2C^{\mathfrak{a}}_{ii}(t)}{\delta}} - \sqrt{\frac{2C^{\mathfrak{r}}_{ii}(t)}{\delta}}$.

Of course, this argument can be made for any choice of dimension i. Hence, a less conservative, yet valid, choice is

$$\gamma^{\mathfrak{a},\mathfrak{r}}(t) := \max_{i=1,\ldots,D} |\mu^{\mathfrak{a}}_i(t) - \mu^{\mathfrak{r}}_i(t)| - \Lambda^{\mathfrak{a},\mathfrak{r}} - \sqrt{\frac{2C^{\mathfrak{a}}_{ii}(t)}{\delta}} - \sqrt{\frac{2C^{\mathfrak{r}}_{ii}(t)}{\delta}}. \tag{3}$$

Notice, this function has the desirable property of being Lipschitz continuous, provided the mean $\mu^{\mathfrak{q}}_i : I \to \mathbb{R}$ and standard deviation functions $\sigma^{\mathfrak{q}}_{ii} = \sqrt{C^{\mathfrak{q}}_{ii}} : I \to \mathbb{R}_+$ are. In particular, it is easy to show $L(\gamma^{\mathfrak{a},\mathfrak{r}}) \leq \max_{i=1,\ldots,D} L(\mu^{\mathfrak{a}}_i) + L(\mu^{\mathfrak{r}}_i) + \sqrt{\frac{2}{\delta}}\big(L(\sigma^{\mathfrak{a}}_{ii}) + L(\sigma^{\mathfrak{r}}_{ii})\big)$ where, as before, $L(f)$ denotes a Lipschitz constant of function f.

For the special case of two dimensions, we can derive a less conservative alternative criterion function based on a tighter two-dimensional Chebyshev-type bound [28]:

Theorem 2.4 (Alternative collision criterion function)**.** *Let spatial dimensionality be $D = 2$. Choosing*

$r^{\mathfrak{q}}_i(t) := \sqrt{\frac{1}{2\delta^a}} \sqrt{C^{\mathfrak{q}}_{ii}(t) + \frac{\sqrt{C^{\mathfrak{q}}_{ii}(t) C^{\mathfrak{q}}_{jj}(t) (C^{\mathfrak{q}}_{ii}(t) C^{\mathfrak{q}}_{jj}(t) - (C^{\mathfrak{q}}_{ij}(t))^2)}}{C^{\mathfrak{q}}_{jj}(t)}}$

($\mathfrak{q} \in \{\mathfrak{a}, \mathfrak{r}\}, i \in \{1, 2\}, j \in \{1, 2\} - \{i\}$) in Eq. 2 yields a valid distribution-independend criterion function. That is, $\gamma^{\mathfrak{a},\mathfrak{r}}(t) > 0 \Rightarrow \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] < \delta^{\mathfrak{a}}$.

A proof sketch and a Lipschitz constant (for non-zero uncertainty) are provided in the appendix. Note, the Lipschitz constant we have derived therein becomes infinite in the limit of vanishing variance. In that case, the presence of negative criterion values can be tested based on the sign of the minimum of the criterion function. This can be found employing a global optimiser. Future work will investigate, in how far Hoelder continuity instead of Lipschitz continuity can be leveraged to yield a similar algorithm as the one provided in Sec. 2.1.2.

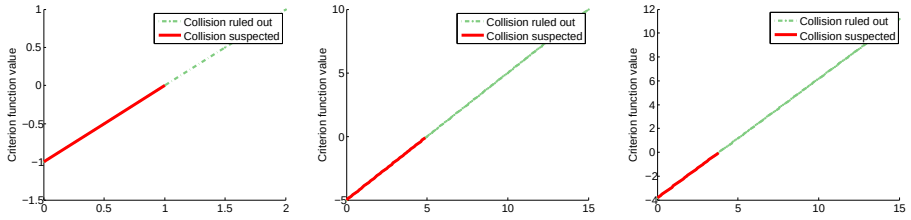

Figure 2: Criterion function values (as per Eq. 3) as a function of $\|\langle x^{\mathfrak{r}}\rangle - \langle x^{\mathfrak{a}}\rangle\|_\infty$ and with $\delta = 0.05$, $\Lambda^{\mathfrak{a},\mathfrak{r}} = 1$. Left: variances $C^{\mathfrak{a}} = C^{\mathfrak{r}} = \mathrm{diag}(.00001, .00001)$. Centre: variances $C^{\mathfrak{a}} = C^{\mathfrak{r}} = \mathrm{diag}(.1, .1)$. Right: variances $C^{\mathfrak{a}} = C^{\mathfrak{r}} = \mathrm{diag}(.1, .1)$ and with improved criterion function (as per Thm. B.3).

2.2.2 Multi-agent case.

Let $\mathfrak{a} \in \mathfrak{A}$, $\mathfrak{A}' \subset \mathfrak{A}$ such that $\mathfrak{a} \notin \mathfrak{A}'$ a subset of agents. We define the event that $\mathfrak{a}$ collides with at least one of the agents in $\mathfrak{A}$' at time t as $\mathfrak{C}^{\mathfrak{a},\mathfrak{A}'}(t) := \{(x^{\mathfrak{a}}(t), x^{\mathfrak{r}}(t)) | \exists \mathfrak{r} \in \mathfrak{A}' : \|x^{\mathfrak{a}}(t) - x^{\mathfrak{r}}(t)\|_2 \leq \Lambda\} = \bigcup_{\mathfrak{r}\in\mathfrak{A}'} \mathfrak{C}^{\mathfrak{a},\mathfrak{r}}$. By union bound, $\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{A}'}(t)] \leq \sum_{\mathfrak{r}\in\mathfrak{A}'} \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)]$.

Theorem 2.5 (Multi-Agent Criterion)**.** *Let $\gamma^{\mathfrak{a},\mathfrak{r}}$ be valid criterion functions defined w.r.t. collision bound $\delta^{\mathfrak{a}}$. We define* multi-agent collision criterion function $\Gamma^{\mathfrak{a},\mathfrak{A}'}(t) := \min_{\mathfrak{r}\in\mathfrak{A}'} \gamma^{\mathfrak{a},\mathfrak{r}}(t)$. *If $\Gamma^{\mathfrak{a},\mathfrak{A}'}(t) > 0$ then the collision probability with $\mathfrak{A}$' is bounded below $\delta^{\mathfrak{a}}|\mathfrak{A}'|$. That is,* $\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{A}'}(t)] < \delta^{\mathfrak{a}}|\mathfrak{A}'|$.

Proof. Let $\mathfrak{a} \in \mathfrak{A}$, $\mathfrak{A}' \subset \mathfrak{A}$ such that $\mathfrak{a} \notin \mathfrak{A}'$ a subset of agents. We define the event that $\mathfrak{a}$ collides with at least one of the agents in $\mathfrak{A}$' at time t as $\mathfrak{C}^{\mathfrak{a},\mathfrak{A}'}(t) := \{(x^{\mathfrak{a}}(t), x^{\mathfrak{r}}(t)) | \exists \mathfrak{r} \in \mathfrak{A}' : \|x^{\mathfrak{a}}(t) - x^{\mathfrak{r}}(t)\|_2 \leq \Delta\} = \bigcup_{\mathfrak{r}\in\mathfrak{A}'} \mathfrak{C}^{\mathfrak{a},\mathfrak{r}}$.

We have established that if $\forall \mathfrak{r} \in \mathfrak{A}' : \gamma^{\mathfrak{a},\mathfrak{r}}(t) > 0$ then $\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] < \delta^{\mathfrak{a}}, \forall \mathfrak{r} \in \mathfrak{A}'$. Now, let $\Gamma^{\mathfrak{a},\mathfrak{A}'} < \delta^{\mathfrak{a}}$. Hence,$\forall \mathfrak{r} \in \mathfrak{A}' : \gamma^{\mathfrak{a},\mathfrak{r}}(t) > 0$. Thus, $\forall \mathfrak{r} \in \mathfrak{A}' : \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] < \delta^{\mathfrak{a}})$ Therefore, $\sum_{\mathfrak{r}\in\mathfrak{A}'} \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] \leq |\mathfrak{A}'|\,\delta^{\mathfrak{a}}$. By union bound, $\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{A}'}(t)] \leq \sum_{\mathfrak{r}\in\mathfrak{A}'} \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)]$. Consequently, we have $\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{A}'}(t)] \leq |\mathfrak{A}'|\,\delta^{\mathfrak{a}}$. q.e.d.

□

Moreover, $\Gamma^{\mathfrak{a},\mathfrak{A}'}$ is Lipschitz if the constituent functions $\gamma^{\mathfrak{a},\mathfrak{r}}$ are (see Appendix B).

Our distribution-independent collision criterion functions have the virtue that they work for all distributions – not only the omnipresent Gaussian. Unfortunately, distribution-independence is gained at the price of conservativeness (ref. to Fig. 2). In our experiments in Sec. 4, the collision criterion function as per Thm. B.3 is utilized as an integral component of our collision avoidance mechanisms. The results suggest that the conservativeness of our detection module does not entail prohibitively high-false-alarm rates for the distribution-independent approach to be considered impractical. That said, whenever distributional knowledge can be converted into a criterion function. One could then

use our derivations as a template to generate refined criterion functions using Eq. 2 with adjusted radii r_i,r_j, reflecting the distribution at hand.

3 Collision Avoidance

In this section we outline the core ideas of our proposed approach to multi-agent collision avoidance. After specifying the agent's dynamics and formalizing the notion of a single-agent plan, we define the multi-agent planning task. Then we describe how conflicts, picked-up by our collision prediction method, can be resolved. In Sec. 3.1 we describe the two coordination approaches we consider utilizing to generate conflict-free plans.

I) Model (example). We assume the system contains a set $\mathfrak{A}$ of agents indexed by $\mathfrak{a} \in \{1, ..., |\mathfrak{A}|\}$. Each agent $\mathfrak{a}$'s associated plant has a probabilistic state trajectory following stochastic controlled D-dimensional state dynamics (we consider the case $D = 2$) in the continuous interval of (future) time $I = (t_0, t_f]$. We desire to ask agents to adjust their policies to avoid collisions. Each policy gives rise to a stochastic belief over the trajectory resulting from executing the policy. For our method to work, all we require is that the trajectory's mean function $m : I \to \mathbb{R}^D$ and covariance matrix function $\Sigma : I \to \mathbb{R}^{D\times D}$ are evaluable for all times $t \in I$.

A prominent class for which closed-form moments can be easily derived are linear stochastic differential equations (*SDE*s). For instance, we consider the SDE

$$dx^{\mathfrak{a}}(t) = K\big(\xi^{\mathfrak{a}}(t) - x^{\mathfrak{a}}(t)\big)dt + B\,dW \tag{4}$$

where $K, B \in \mathbb{R}^{D\times D}$ are matrices $x^{\mathfrak{a}} : I \to \mathbb{R}^D$ is the state trajectory and W is a vector-valued Wiener process. Here, $u(x^{\mathfrak{a}}; \xi^{\mathfrak{a}}) := K(\xi^{\mathfrak{a}} - x^{\mathfrak{a}})$ could be interpreted as the control policy of a linear feedback-controller parametrised by $\xi^{\mathfrak{a}}$. It regulates the state to track a desired trajectory $\xi^{\mathfrak{a}}(t) = \zeta_0^{\mathfrak{a}}\chi_{\{0\}}(t) + \sum_{i=1}^{H^{\mathfrak{a}}} \zeta_i^{\mathfrak{a}}\chi_{\tau_i^{\mathfrak{a}}}(t)$ where $\chi_{\tau_i} : \mathbb{R} \to \{0, 1\}$ denotes the indicator function of the half-open interval $\tau_i^{\mathfrak{a}} = (t_{i-1}^{\mathfrak{a}}, t_i^{\mathfrak{a}}] \subset [0, T^{\mathfrak{a}}]$ and each $\zeta_i^{\mathfrak{a}} \in \mathbb{R}^D$ is a *setpoint*. If K is positive definite the agent's state trajectory is determined by setpoint sequence $p^{\mathfrak{a}} = (t_i^a, \zeta_i^{\mathfrak{a}})_{i=0}^{H^{\mathfrak{a}}}$ (aside from the random disturbances) which we will refer to as the agent's *plan*. For example, plan $p^{\mathfrak{a}} := \big((t_0, x_0^{\mathfrak{a}}), (t_f, x_f^{\mathfrak{a}})\big)$ could be used to regulate agent $\mathfrak{a}$'s *start state* $x_0^{\mathfrak{a}}$ to a given *goal state* $x_f^{\mathfrak{a}}$ between times t_0 and t_f. For simplicity, we assume the agents are always initialized with plans of this form before coordination commences.

One may interpret a setpoint as some way to alter the stochastic trajectory. Below, we will determine setpoints that modify a stochastic trajectory to reduce collision probability while maintaining low expected cost. From the vantage point of policy search, $\xi^{\mathfrak{a}}$ is agent $\mathfrak{a}$'s policy parameter that has to be adjusted to avoid collisions.

II) Task. Each agent $\mathfrak{a}$ desires to find a sequence of setpoints $(p^{\mathfrak{a}})$ such that (i) it moves from its start state $x_0^{\mathfrak{a}}$ to its goal state $x_f^{\mathfrak{a}}$ along a low-cost trajectory and (ii) such that along the trajectory its plant (with diameter Δ)

does not collide with any other agents' plant in state space with at least a given probability $1 - \delta \in (0, 1)$.

III) Collision resolution. An agent seeks to avoid collisions by adding new setpoints to its plan until the collision probability of the resulting state trajectory drops below threshold δ. For choosing these new setpoints we consider two methods **WAIT** and **FREE**. In the first method the agents insert a time-setpoint pair $(t, x_0^{\mathfrak{a}})$ into the previous plan $p^{\mathfrak{a}}$. Since this aims to cause the agent to wait at its start location $x_0^{\mathfrak{a}}$ we will call the method WAIT. It is possible that multiple such insertions are necessary until collisions are avoided. Of course, if a higher-priority agent decides to traverse through $x_0^{\mathfrak{a}}$, this method is too rigid to resolve a conflict. In the second method the agent optimizes for the time and location of the new setpoint. Let $p^{\mathfrak{a}}_{\uparrow(t,s)}$ be the plan updated by insertion of time-setpoint pair $(t, s) \in I \times \mathbb{R}^D$. We propose to choose the candidate setpoint (t, s) that minimizes a function being a weighted sum of the expected cost entailed by executing updated plan $p^{\mathfrak{a}}_{\uparrow(t,s)}$ and a hinge-loss collision penalty $c^{\mathfrak{a}}_{coll}(p^{\mathfrak{a}}_{\uparrow(t,s)}) := \lambda \max\{0, -\min_t \Gamma^{\mathfrak{a}}(t)\}$. Here, $\Gamma^{\mathfrak{a}}$ is computed based on the assumption we were to execute $p^{\mathfrak{a}}_{\uparrow(t,s)}$ and $\lambda >> 0$ determines the extent to which collisions are penalized. Since the new setpoint can be chosen freely in time and state-space we refer to the method as FREE.

3.1 Coordination

We will now consider how to integrate our collision detection and avoidance methods into a coordination framework that determines who needs to avoid whom and at what stage of the coordination process. Such decisions are known to significantly impact the *social cost* (i.e. the sum of all agents' individual costs) of the agent collective.

Fixed-priorities (FP). As a baseline method for coordination we consider a basic fixed-priority method (e.g. [3,11]). Here, each agent has a unique ranking (or priority) according to its index $\mathfrak{a}$ (i.e. agent 1 has highest priority, agent $|\mathfrak{A}|$ lowest). When all higher-ranking agents are done planning, agent $\mathfrak{a}$ is informed of their planned trajectories which it has to avoid with a probability greater than $1 - \delta$. This can be done by repeatedly invoking for collision detection and resolution methods described above until no further collision with higher-ranking agents are found.

Lazy Auction Protocol (AUC). While the FP method is simple and fast the rigidity of the fixed ranking can lead to sub-optimal social cost and coordination success. Furthermore, its sequential nature does not take advantage of possible parallelization a distributed method could. To alleviate this we propose to revert the ranking flexibly on a case-by-case basis. In particular, the agents are allowed to compete for the right to gain passage (e.g. across a region where a collision was detected) by submitting bids in the course of an auction. The structure of the approach is outlined in Alg. 3.

Assume an agent $\mathfrak{a}$ detects a collision at a particular time step t_{coll} and

input : Agents $\mathfrak{a} \in \mathfrak{A}$, cost functions $c^{\mathfrak{a}}$, dynamics, initial start and goal states, initial plans $p^1, ..., p^{|\mathfrak{A}|}$.
output: collision-free plans $p^1, ..., p^{|\mathfrak{A}|}$.

repeat
  for $\mathfrak{a} \in \mathfrak{A}$ **do**
    $[\ \mathsf{flag}^{\mathfrak{a}}, \mathcal{C}^{\mathfrak{a}}, t_{\mathsf{coll}}] \leftarrow \texttt{CollDetect}^{\mathfrak{a}}(\mathfrak{a}, \mathfrak{A} - \{\mathfrak{a}\})$
    if $\mathsf{flag}^{\mathfrak{a}} = 1$ **then**
      $\mathsf{winner} \leftarrow \texttt{Auction}(\mathcal{C}^{\mathfrak{a}} \cup \{\mathfrak{a}\}, t_{\mathsf{coll}})$
      foreach $\mathfrak{r} \in (\mathcal{C}^{\mathfrak{a}} \cup \{\mathfrak{a}\}) - \{\mathsf{winner}\}$ **do**
        $p^{\mathfrak{r}} \leftarrow \texttt{Avoid}^{\mathfrak{r}}((\mathcal{C}^{\mathfrak{a}} \cup \{\mathfrak{a}\}) - \{\mathfrak{r}\}, t_{\mathsf{coll}})$
        $\texttt{Broadcast}^{\mathfrak{r}}\ (p^{\mathfrak{r}})$
      end
    end
  end
until $\forall \mathfrak{a} \in \mathfrak{A} : \mathsf{flag}^{\mathfrak{a}} = 0$;

Algorithm 3: Lazy auction coordination method (AUC) (written in a sequentialized form). Collisions are resolved by choosing new setpoints to enforce collision avoidance. $\mathcal{C}^{\mathfrak{a}}$: set of agents detected to be in conflict with agent $\mathfrak{a}$. $\mathsf{flag}^{\mathfrak{a}}$: collision detection flag (=0, iff no collision detected). t_{coll}: earliest time where a collision was detected. Avoid: collision resolution method updating the plan by a single new setpoint according to WAIT or FREE.

invites the set of agents $\mathcal{C}^{\mathfrak{a}} = \{\mathfrak{r} | \gamma^{\mathfrak{a},\mathfrak{r}}(t_{\mathrm{coll}}) \leq 0\}$ to join an auction to decide who needs to avoid whom. In particular, the auction determines a winner who is not required to alter his plan. The losing agents need to insert a new setpoint into their respective plans designed to avoid all other agents in $\mathcal{C}^{\mathfrak{a}}$ while keeping the plan cost function low.

The idea is to design the auction rules as a heuristic method to minimize the social cost of the ensuing solution. To this end, we define the bids such that their magnitude is proportional to a heuristic magnitude of the expected regret for losing and not gaining passage. That is agent $\mathfrak{a}$ submits a bid $b^{\mathfrak{a}} = \mathfrak{l}^{\mathfrak{a}} - \mathfrak{s}^{\mathfrak{a}}$. Magnitude $\mathfrak{l}^{\mathfrak{a}}$ is defined as $\mathfrak{a}$'s anticipated cost $c^{\mathfrak{a}}_{plan}(p^{\mathfrak{a}}_{\uparrow(t,s)})$ for the event that the agent will not secure "the right of passage" and has to create a new setpoint (t, s) (according to (III)) tailored to avoid all other agents engaged in the current auction. On the other hand, $\mathfrak{s}^{\mathfrak{a}} := c^{\mathfrak{a}}_{plan}(p^{\mathfrak{a}})$ is the cost of the unchanged plan $p^{\mathfrak{a}}$. If there is a tie among multiple agents the agent with the lowest index among the highest bidders wins.

Acknowledging that $\mathfrak{s}^{winner} + \sum_{\mathfrak{a} \neq winner} \mathfrak{l}^{\mathfrak{a}}$ is an estimated social cost (based on current beliefs of trajectories) after the auction, we see that the winner determination rule greedily attempts to minimize social cost: $b^{winner} \geq b^{\mathfrak{r}} \Leftrightarrow \forall \mathfrak{r} : \mathfrak{s}^{\mathfrak{r}} + \sum_{\mathfrak{a} \neq \mathfrak{r}} \mathfrak{l}^{\mathfrak{a}} \geq \mathfrak{s}^{winner} + \sum_{\mathfrak{a} \neq winner} \mathfrak{l}^{\mathfrak{a}}$.

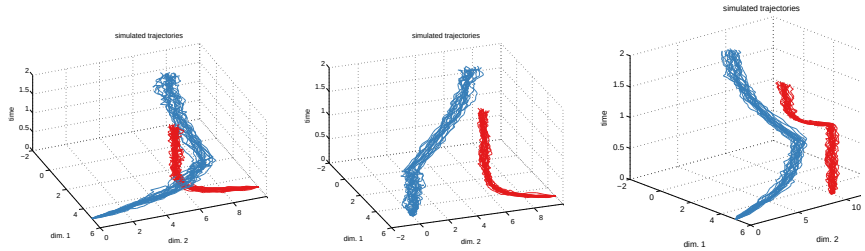

Figure 3: EXP1. Draws from uncoordinated agents' plans (left), after coordination and collision resolution with methods FP-WAIT (centre) and AUC-WAIT (right).

	Experiment 1			Experiment 2		
Quantity	NONE	AUC-WAIT	FP-WAIT	NONE	AUC-FREE	FP-FREE
A	78	0	0	51	0	0
B	13.15	13.57	12.57	14.94	16.22	18.13
C	0.05	0.04	25.8	0.05	0.05	0.05
D	0	6	3	0	4	4

Table 1: Quantities estimated based on 100 draws from SDEs simulating executions of the different plans in EXP1 and EXP2. **A**: estimated collision probability [%]; **B**: averaged path length away from goal; **C**: averaged sqr. dist. of final state to goal; **D**: number of collision resolution rounds. Notice, our collision avoidance methods succeed in preventing collisions. In EXP1 the FP-WAIT method failed to reach its first goal in time which is reflected in the *sqr. distance to goal* measure. Note the discrepancies in avg. path length are relatively low due to convexity effects and the contribution of state noise to the path lengths.

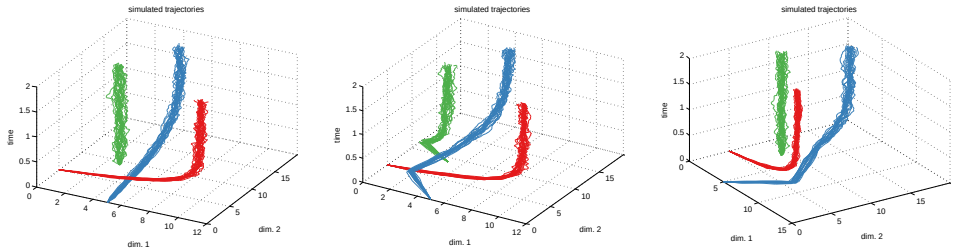

Figure 4: EXP2. Draws from uncoordinated agents' plans (left), after coordination and collision resolution with methods FP-FREE (centre) and AUC-FREE (right).

4 Simulations

As a first test, we simulated three simple multi-agent scenarios, *EXP1*, *EXP2* and *EXP3*. Each agent's dynamics were an instantiation of an SDE of the form of Eq. 4. We set δ to achieve collision avoidance with certainty greater than 95%. Collision prediction was based on the improved criterion function as per Thm. B.3. During collision resolution with the FREE method each agent $\mathfrak{a}$ assessed a candidate plan $p^{\mathfrak{a}}$ according to cost function $c^{\mathfrak{a}}_{plan}(p^{\mathfrak{a}}) = w_1\, c^{\mathfrak{a}}_{traj}(p^{\mathfrak{a}}) + w_2\, c^{\mathfrak{a}}_{miss}(p^{\mathfrak{a}}) + w_3\, c^{\mathfrak{a}}_{coll}(p^{\mathfrak{a}})$. Here $c^{\mathfrak{a}}_{traj}$ is a heuristic to penalize expected control energy or path length; in the second summand, $c^{\mathfrak{a}}_{miss}(p^{\mathfrak{a}}) = \left\| x^{\mathfrak{a}}(t_f) - x^{\mathfrak{a}}_f \right\|^2$ penalizes expected deviation from the goal state; the third term $c^{\mathfrak{a}}_{coll}(p^{\mathfrak{a}})$ penalizes collisions (cf. III). The weights are design parameters which we set to $w_1 = 10, w_2 = 10^3$ and $w_3 = 10^6$, emphasizing avoidance of mission failure and collisions. Note, if our method was to be deployed in a receding horizon fashion, the parameters could also be adapted online using standard learning techniques such as no-regret algorithms [18, 24].

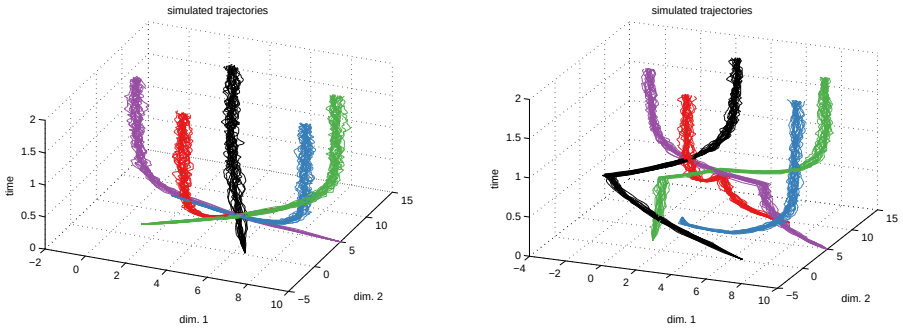

Figure 5: Ex. of EXP3 with 5 agents. Draws from uncoordinated agents' plans (left), after coordination and collision resolution with methods AUC-FREE (right).

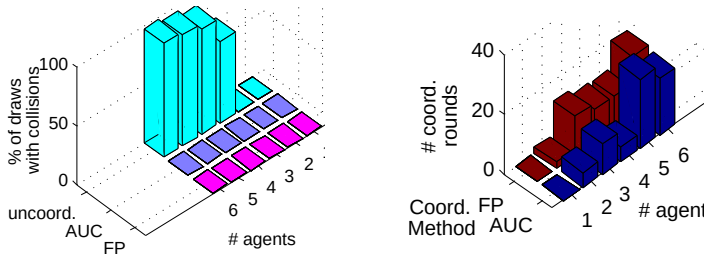

Figure 6: Recorded results for EXP3 with 1 to 6 agents. Note, all collisions were successfully avoided.

EXP1. Collision resolution was done with the WAIT method to update

plans. Draws from the SDEs with the initial plans of the agents are depicted in Fig. 3 (left). The curves represent 20 noisy trajectories of agents 1 (red) and 2 (blue). Each curve is a draw from the stochastic differential dynamics obtained by simulating the execution of the given initial plan. The trajectories were simulated with the Euler-Maruyama method for a time interval of $I = [0s, 2s]$. The spread of the families of curves is due to the random disturbances each agent's controller had to compensate for during runtime.

Agent 1 desired to control the state from start state $x_0^1 = (5, 10)$ to goal $x_f^1 = (5, 5)$. Agent 2 desired to move from start state $x_0^2 = (5, 0)$ via intermediate goal $x_{f_1}^2 = (5, 7)$ (at 1s) to final goal state $x_{f_2}^2 = (0, 7)$. While the agents meet their goals under the initial plans, their execution would imply a high probability of colliding around state $(5, 6)$ (cf. Fig. 3 (left), Tab. 1). Coordination with fixed priorities (1 (red) $>$ 2 (blue)) yields conflict-free plans (Fig. 3 (centre)). However, agent 2 is forced to wait too long at its start location to be able to reach intermediate waypoint $x_{f,1}^2$ in time and therefore, decides to move directly to its second goal. This could spawn high social cost due to missing one of the designated goals (Tab. 1). By contrast, the auction method is flexible enough to reverse the ranking at the detected collision point causing agent 1 to wait instead of 2 (Fig. 3 (right)). Thereby, agent 2 is able to reach both of its goal states in time. This success is reflected by low social cost (see Tab. 1).

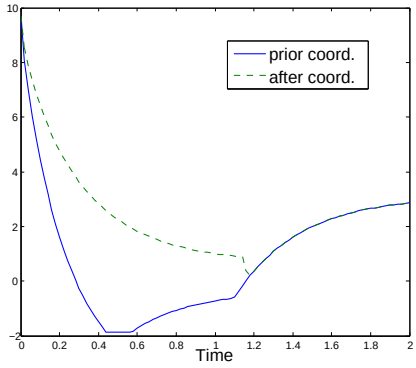

Figure 7: EXP1. Criterion functions for collision detection of agent 2 before and after coordination. The first graph accurately warns of a collision before coordination (as indicated by negative values), whereas the one corresponding to collision-free trajectories is in the positive-half space as desired.

EXP2. The setup was analogous to EXP1 but with three agents and different start and goal states as depicted in Fig. 4. Furthermore, collisions were avoided with the FREE method with 10 random initializations of the local optimizer. Coordination of plans with fixed priorities (1 (red) $>$ 2 (blue) $>$ 3 (green)) caused 2 to avoid agent 1 by moving to the left. Consequently, 3 now had to temporarily leave its start and goal state to get out of the way (see Fig. 4 (centre)). With two agents moving to avoid collisions social cost was relatively

high (see Tab. 1). During coordination with the auction-based method agent 2 first chose to avoid agent 1 (as in the FP method). However, losing the auction to agent 3 at a later stage of coordination, agent 2 decided to finally circumvent 1 by arcing to the right instead of to the left. This allowed 3 to stay in place (see Tab. 1).

EXP3. Next, we conducted a sequence of experiments for varying numbers of agents ranging from $|\mathfrak{A}| = 1, .., 7$. In each experiment all agents' start locations were placed on a circle. Their respective goals were placed on the opposite ends of the circle. The eigenvalues of the feedback gain matrices of each agent were drawn at random from a uniform distribution on the range [2,7]. An example situation for an experiment with 5 agents is depicted in Fig. 5. Collision avoidance was achieved.

Note, that despite this setting being close to worst case (i.e. almost all agents try to traverse a common, narrow corridor) the coordination overhead is moderate (see Fig. 6, right). Also, all collisions were successfully avoided (see Fig. 6, left).

5 Conclusions

This work considered multi-agent planning under stochastic uncertainty and non-convex chance-constraints for collision avoidance. In contrast to pre-existing work, we did not need to rely on prior space or time-discretisation. This was achieved by deriving criterion functions with the property that the collision probability is guaranteed to be below a freely definable threshold $\delta \in (0, 1)$ if the criterion function attains no negative values. Thereby, stochastic collision detection is reduced to deciding whether such negative values exist. For Lipschitz criterion functions, we provided an algorithm for making this decision rapidly. We described a general procedure for deriving criterion functions and presented two such functions based on Chebyshev-type bounds. The advantage of using Chebyshev inequalities is their independence of the underlying distribution. Therefore, our approach is applicable to any stochastic state noise model for which the first two moments can be computed at arbitrary time steps. In particular, this would apply to models with state-dependent uncertainty and non-convex chance constraints which, to the best to our knowledge, have not been successfully approached in the multi-agent control literature. Nonetheless, future work could build on our results and derive less conservative criterion functions by using more problem-specific probabilistic inequalities. For instance, in simple cases such as additive Gaussian noise, tighter bounds can be given [4] and used in Eq. 2.

To enforce collision avoidance, our method modified the agent's plans until no collisions could be detected. To coordinate the detection and avoidance efforts of the agents, we employed an auction-based as well as a fixed-priority method.

Our experiments are a first indication that our approach can succeed in

finding collision-free plans with high-certainty with the number of required coordination rounds scaling mildly in the number of agents. While in its present form, the coordination mechanism does not come with a termination guarantee, in none of our simulations have we encountered an infinite loop. For graph routing, [7] provides a termination guarantee of the lazy auction approach under mild assumptions. Current work considers if their analysis can be extended to our continuous setting. Moreover, if required, our approach can be combined with a simple stopping criterion that terminates the coordination attempt when a computational budget is expended or an infinite loop is detected.

The computation time within each coordination round depends heavily on the time required for finding a new setpoint and for collision detection. This involves minimizing $(t, s) \mapsto c^{\mathfrak{a}}_{plan}(p^{\mathfrak{a}}_{\uparrow(t,s)})$ and $c^{\mathfrak{a}}_{coll}$, respectively. The worst-case complexity depends on the choice of cost functions, their domains and the chosen optimizer. Fortunately, we can draw on a plethora of highly advanced global optimisation methods (eg [13, 23]) guaranteeing rapid optimization success. In terms of execution time, we can expect considerable alleviations from implementation in a compiled language. Furthermore, the collision detection and avoidance methods are based on global optimization and thus, would be highly amenable to parallel processing – this could especially benefit the auction approach.

While our exposition was focussed on the task of defining setpoints of feedback-controlled agents, the developed methods can be readily applied to other policy search settings, where the first two moments of the probabilistic beliefs over the trajectories (that would result from applying the found policies) can be computed.

References

[1] N. Ayanian and V. Kumar. Decentralized feedback controllers for multiagent teams in environments with obstacles. *IEEE Trans. on Robotics*, 26(5), 2010.

[2] D. Bareiss and J. van den Berg. Reciprocal collision avoidance for quadrotor helicopters using lqr-obstacles. In *AAAI-12 Workshop on Multiagent Pathfinding*, 2012.

[3] M. Bennewitz, W. Burgard, and S. Thrun. Exploiting constraints during prioritized path planning for teams of mobile robots. In *IROS*, 2001.

[4] L. Blackmore, H. Li, and B. Williams. A probabilistic approach to optimal robust path planning with obstacles. In *American Control Conference*, 2006.

[5] J. Calliess, M. Osborne, and S. Roberts. Towards auction-based multi-agent collision-avoidance under continuous stochastic dynamics. Presented at *ICML-2012, Workshop on Markets, Mechanisms, and Multi-Agent Models - Examining the Interaction of Machine Learning and Economics*, 2012.

[6] J. Calliess, M. Osborne, and S. Roberts. Conservative collision prediction and avoidance for stochastic trajectories in continuous time and space. In *Proc. of the 13th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, 2014.

[7] J. Calliess, D. Lyons, and U. Hanebeck. Lazy auctions for multi-robot collision avoidance and motion control under uncertainty. Technical Report PARG-01-11, Dept. of Engineering Science, University of Oxford, 2011.

[8] D.E. Chang, S.C. Shadden, J.E. Marsden, and R. Olfati-Saber. Collision avoidance for multiple agent systems. In *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, volume 1, pages 539–543. IEEE, 2003.

[9] D.V. Dimarogonas, S.G. Loizou, K.J. Kyriakopoulos, and M.M. Zavlanos. A feedback stabilization and collision avoidance scheme for multiple independent non-point agents. *Automatica*, 42(2):229–243, 2006.

[10] M. G. Earl and R. D'Andrea. Iterative MILP Methods for Vehicle-Control Problems. *IEEE Trans. on Robotics*, 2005.

[11] M. A. Erdmann and T. Lozano-Perez. On multiple moving objects. *Algorithmica*, 1987.

[12] C. Gardiner. *Stochastic Methods-A Handbook for the Natural and Social Sciences*. Springer, 4th edition, 2009.

[13] D.R. Jones, C.D. Peritunen, and B.E. Stuckman. Lipschitzian optimization without the lipschitz constant. *JOTA*, 79(1), 1991.

[14] Sven Koenig, Craig Tovey, X. Zheng, and I. Sungur. Sequential bundle-bid single-sale auction algorithms for decentralized control. In *IJCAI*, 2007.

[15] D. Kostic, S. Adinandra, J. Caarls, and H. Nijmeijer. Collision-free motion coordination of unicycle multi-agent systems. In *American Control Conference (ACC), 2010*, pages 3186–3191. IEEE, 2010.

[16] M. Lagoudakis, V. Markakis, D. Kempe, P. Keskinocak, S. Koenig, A. Kleywegt, C. Tovey, A. Meyerson, and S. Jain. Auction-based multi-robot routing. In *Int. Conf. on Robotics: Science and Systems*, 2005.

[17] Y. Li and K. Gupta. Motion planning of multiple agents in virtual environments on parallel architectures. In *ICRA*, 2007.

[18] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. In *IEEE Symposium on Foundations of Computer Science*, pages 256–261, 1989.

[19] D. Lyons, J. Calliess, and U. Hanebeck. Chance constrained model predictive control for multi-agent systems with coupling constraints. In *American Control Conference (ACC)*, 2012.

[20] Daniel Lyons, J.P. Calliess, and U.D. Hanebeck. Chance-constrained Model Predictive Control for Multi-Agent Systems. *Arxiv preprint arXiv:1104.5384*, 2011.

[21] S. Mastellone, D.M. Stipanović, C.R. Graunke, K.A. Intlekofer, and M.W. Spong. Formation control and collision avoidance for multi-agent non-holonomic systems: Theory and experiments. *The International Journal of Robotics Research*, 27(1):107–126, 2008.

[22] T. Sandholm. Algorithm for optimal winner determination on combinatorial auctions. *Artif. Int.*, 2002.

[23] B. Shubert. A sequential method seeking the global maximum of a function. *SIAM J. on Numerical Analysis*, 9, 1972.

[24] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian Process Optimization in the Bandit Setting : No Regret and Experimental Design. In *ICML*, 2010.

[25] D.M. Stipanović, P.F. Hokayem, M.W. Spong, D.D. Šiljak, et al. Cooperative avoidance control for multiagent systems. *Journal of Dynamic Systems, Measurement, and Control*, 129:699, 2007.

[26] C. Tovey, M. Lagoudakis, S. Jain, and S. Koenig. The generation of bidding rules for auction-based robot coordination. In *Multi-Robot Systems: From Swarms to Intelligent Automata*, volume 3, pages 3–14, 2005.

[27] J. Van den Berg, M. Lin, and D. Manocha. Reciprocal velocity obstacles for real-time multi-agent navigation. In *ICRA*, 2008.

[28] P. Whittle. A multivariate generalization of Tchebichev's inequality. *Quarterly Journal of Mathem.*, 1958.

Supplementary derivations.

A Derivations of the covariance and mean of the feedback-controlled agents as deployed in the trajectory planning experiments

We will now solve the mean and covariance for the dynamics given by the Ito-SDE describing the dynamics of the feedback controlled agents considered in the experiments of this paper. Our aim is to obtain closed-form solutions avoiding the need to approximate any integrals. This ameliorates the computational burden of our method since mean and covariance matrix functions need to be evaluated frequently in the course of collision detection and resolution.

Theorem A.1. *For all $t \in [t_0, T]$ let $\xi(t), x(t) \in \mathbb{R}^D$, $A = diag(a_1, ..., a_D), K = diag(k_1, ..., k_D)$, $B = diag(\sqrt{\nu}_1, ..., \sqrt{\nu}_D)$ and let $x(t_0)$ be a normally distributed random vector. The solution to the SDE $\mathbf{dx} = (\mathbf{Ax} - \mathbf{K}(\mathbf{x} - \xi))\mathbf{dt} + \mathbf{B}\,\mathbf{dW}$ is a Gaussian process with vector-valued mean function $\mu : [t_0, T] \to \mathbb{R}^D$ and matrix-valued covariance function $C : [t_0, T]^2 \to \mathbb{R}^{D\times D}$. Here the mean components are*

$$\mu_j(t) = e^{(k_j - a_j)(t_0 - t)} \langle x_j(t_0) \rangle + \int_{t_0}^{t} k_j\, e^{(k_j - a_j)(\tilde{t} - t)} \xi_j(\tilde{t}) d\tilde{t}$$

and the covariance matrix function is $C(s,t) = diag(cov_{11}(s,t), ..., cov_{DD}(s,t))$ where

$$\begin{aligned} cov_{jj}(s,t) = {} & e^{(a_j - k_j)(t + s - 2t_0)} (\langle x_j^2(t_0) \rangle - \langle x_j(t_0) \rangle^2) \\ & + \frac{\nu_j}{2(k_j - a_j)} [e^{(a_j - k_j)|t-s|} - e^{(k_j - a_j)(2t_0 - (s+t))}]. \end{aligned}$$

Proof. Let $q_j := k_j - a_j$. Owing to the diagonal form of A, K and B and the independence of the output dimensions of the vector-valued Wiener process, the given SDE decomposes into a system of D indpendent 1-dimensional SDEs $dx_j = (-q_j x_j + k_j \xi_j)dt + \sqrt{\nu}_j dW_j$ $(j = 1, ..., D)$ which can be treated separately.

For ease of notation we omit the subscripts yielding an SDE of the form: $dx = (-qx + k\xi)dt + \sqrt{\nu}\, dW$. To solve each of these SDEs we introduce the substitution $y := xe^{qt}$. With Ito's product rule we find

$$\begin{aligned} dy &= x\, de^{qt} + e^{qt}\, dx + dx\, de^{qt} \\ &= \xi k e^{qt}\, dt + \sqrt{\nu} e^{qt} dW \end{aligned}$$

where the last equality follows from substitution of the SDE for dx and utilization the formal rules $dt\, dW = 0, (dt)^2 = 0$.

This SDE is solved by $y(t) = y(t_0) + \int_{t_0}^{t} k\xi(\tilde{t}) e^{q\tilde{t}} d\tilde{t} + \int_{t_0}^{t} e^{q\tilde{t}} \sqrt{\nu} dW(\tilde{t})$ where the last integral is to be interpreted as an Ito integral. Re-substitution yields

$$x(t) = x(t_0) e^{q(t_0 - t)} + \int_{t_0}^{t} k\, e^{q(\tilde{t} - t)} \xi(\tilde{t}) d\tilde{t} + \int_{t_0}^{t} e^{q(\tilde{t} - t)} \sqrt{\nu} dW(\tilde{t}).$$

Random variables $x(t)$ ($t \in [t_0, T]$) are given as affine transformations of a Wiener process. Since the Wiener process is Gaussian, the solution is clearly a Gaussian process as long as $x(t_0)$ is a normally distributed random variable.

For notational convenience we define

$$\begin{aligned} J_t &:= \int_{t_0}^{t} k\, e^{q(\tilde{t}-t)} \xi(\tilde{t})\, d\tilde{t}, \\ \Omega_t &:= \int_{t_0}^{t} e^{q(\tilde{t}-t)} \sqrt{\nu}\, dW(\tilde{t}) \\ \text{and } c_t &:= x(t_0) e^{q(t_0-t)}. \end{aligned}$$

Therefore, we can write

$$x(t) = c_t + J_t + \Omega_t.$$

It remains to find the first and second moments. Owing to the linearity of the expectation and, since the expectation of the Ito integral of a non-anticipating integrand is zero, the mean function is:

$$\mu(t) = \langle x(t) \rangle = \langle c_t \rangle + \langle J_t \rangle + \langle \Omega_t \rangle = \langle x(t_0) \rangle e^{k(t_0-t)} + J_t.$$

Now, we calculate the second moment:

$$\operatorname{cov}(s,t) = \langle x(s)x(t) \rangle - \mu(s)\mu(t)$$

where

$$\begin{aligned} \langle x(s)x(t) \rangle &= \langle (c_s + J_s + \Omega_s)(c_t + J_t + \Omega_t) \rangle \\ &= \langle c_s c_t \rangle + \langle c_s J_t \rangle + \langle c_s \Omega_t \rangle + ... + \langle \Omega_s \Omega_t \rangle. \end{aligned}$$

Once again leveraging that the expectation of an Ito integral with non-anticipating integrand is zero, the cross-terms $\langle c_q \Omega_r \rangle, \langle J_q \Omega_r \rangle$ ($q, r \in \{s, t\}$) vanish. Hence, we obtain

$$\begin{aligned} \langle x(s)x(t) \rangle &= \langle c_s c_t \rangle + \langle c_s J_t \rangle + \langle c_t J_s \rangle + J_t J_s + \langle \Omega_s \Omega_t \rangle \\ &= \langle c_s c_t \rangle + J_t \langle c_s \rangle + J_s \langle c_t \rangle + J_t J_s + \langle \Omega_s \Omega_t \rangle \\ &= \mu(s)\mu(t) + \langle c_s c_t \rangle - \langle c_s \rangle \langle c_t \rangle + \langle \Omega_s \Omega_t \rangle. \end{aligned}$$

Hence,

$$\begin{aligned} \operatorname{cov}(s,t) &= \langle c_s c_t \rangle - \langle c_s \rangle \langle c_t \rangle + \langle \Omega_s \Omega_t \rangle \\ &= e^{-q(t+s-2t_0)} (\langle x^2(t_0) \rangle - \langle x(t_0) \rangle^2) + \langle \Omega_s \Omega_t \rangle \end{aligned}$$

It is well-known that for non-anticipating f, g interval I we have

$$\langle \int_I f(t) dW(t) \int_I g(t) dW(t) \rangle = \int_I f(t) g(t) dt.$$

Applying this fact as well as leveraging the independent increments property of the Wiener process yields:

$$\begin{aligned}\langle \Omega_s \Omega_t \rangle &= \frac{\nu}{e^{q(s+t)}} \int_{t_0}^{\min\{t,s\}} e^{2q\tilde{t}} d\tilde{t} \\ &= \frac{\nu}{2qe^{q(s+t)}} [e^{2k\min\{t,s\}} - e^{2qt_0}] \\ &= \frac{\nu}{2q} [e^{q(2\min\{t,s\}-(s+t)} - e^{q(2t_0-(s+t))}] \\ &= \frac{\nu}{2q} [e^{-q|t-s|} - e^{q(2t_0-(s+t))}]\end{aligned}$$

□

Notice, that when altering the plan and hence, ξ, the integral

$$J_{t,j} := \int_{t_0}^{t} k_j \, e^{(k_j - a_j)(\tilde{t}-t)} \xi_j(\tilde{t}) d\tilde{t}$$

has to be computed. For general forms of allowable ξ we would have to rely on numerical approximation methods. Since repeated calculations of solutions need to be done in the course of coordination we will want to alleviate the computational burden thereof as much as possible. This motivated our restriction to plans that give rise to step functions for which we can show that J_t is closed-form.

Lemma A.2. *Let $t \geq t_0, t, t_0 \in [0, T^{\mathfrak{a}}]$. Given plan $p^{\mathfrak{a}} = \big((t_i^{\mathfrak{a}}, \zeta_i^{\mathfrak{a}})\big)_{i=0}^{H^{\mathfrak{a}}}$ where each $\zeta_i^{\mathfrak{a}} = (\zeta_{i,j}^{\mathfrak{a}})_{j=1}^{D}$. Let $\underline{i} = \arg\max_i\{t_i \leq t_0\}$, $\bar{i} = \arg\min_i\{t_i \geq t\}$ and $\mathcal{I} := \{i \in \{1, ..., H^{\mathfrak{a}}\} | \underline{i} < i \leq \bar{i}\}$. Furthermore, let $\xi_j^{\mathfrak{a}}$ denote the jth component of step-function setpoint signal $\xi^{\mathfrak{a}}$. We have*

$$\begin{aligned}&\int_{t_0}^{t} k_j \, e^{q_j(\tilde{t}-t)} \xi_j^{\mathfrak{a}}(\tilde{t}) d\tilde{t} \\ &= \sum_{i \in \mathcal{I}} \frac{k_j}{q_j} \zeta_{i,j}^{\mathfrak{a}} (e^{q_j(\min\{t_i,t\}-t)} - e^{q_j(\max\{t_{i-1},t_0\}-t)}).\end{aligned}$$

Proof. Remember, $\tau_i^{\mathfrak{a}} = (t_{i-1}^{\mathfrak{a}}, t_i^{\mathfrak{a}}] \subset [0, T^{\mathfrak{a}}]$. We have

$$\begin{aligned}&\int_{t_0}^{t} k_j \, \exp(q_j(\tilde{t}-t)) \xi_j^{\mathfrak{a}}(\tilde{t}) d\tilde{t} \\ &= \int_{t_0}^{t} k_j \, \exp(q_j(\tilde{t}-t)) \sum_{i=1}^{H^{\mathfrak{a}}} \zeta_{i,j}^{\mathfrak{a}} \chi_{\tau_i^{\mathfrak{a}}}(\tilde{t}) d\tilde{t} \\ &= \sum_{i \in \mathcal{I}} \int_{\max\{t_0, t_{i-1}\}}^{\min\{t,t_i\}} k_j \, \exp(q_j(\tilde{t}-t)) \zeta_{i,j}^{\mathfrak{a}} d\tilde{t}\end{aligned}$$

Calculation of the anti-derivate and substitution of the integration bounds yields the desired result.

□

We immediately get the following corollary

Corollary A.3. *Let $\mathfrak{a} \in \mathfrak{A}$ be an agent with controlled plant dynamics*

$$dx^{\mathfrak{a}} = (Ax^{\mathfrak{a}} - K(x^{\mathfrak{a}} - \xi^{\mathfrak{a}}))dt + B\,dW$$

where for all $t \in [t_0, T]$: $\xi^{\mathfrak{a}}(t), x^{\mathfrak{a}}(t) \in \mathbb{R}^D$, $A = diag(a_1, ..., a_D)$,

$K = diag(k_1, ..., k_D)$, $B = diag(\sqrt{\nu}_1, ..., \sqrt{\nu}_D)$. Let $x^{\mathfrak{a}}(t_0)$ be a normally distributed random vector. Assume $\mathfrak{a}$'s plan is $p^{\mathfrak{a}} = \left((t_i^{\mathfrak{a}}, \zeta_i^{\mathfrak{a}})\right)_{0=1}^{H^{\mathfrak{a}}}$ where each $\zeta_i^{\mathfrak{a}} = (\zeta_{i,j}^{\mathfrak{a}})_{j=1}^{D}$. Let $t \geq t_0$, $\underline{i} = \arg\max_i\{t_i \leq t_0\}$, $\bar{i} = \arg\min_i\{t_i \geq t\}$ and $\mathcal{I} := \{i \in \{1, ..., H^{\mathfrak{a}}\} | \underline{i} < i \leq \bar{i}\}$. Furthermore, let $\xi_j^{\mathfrak{a}}$ denote the jth component of step-function reference signal $\xi^{\mathfrak{a}}$.

The solution to agent $\mathfrak{a}$'s SDE is a Gaussian process with vector-valued mean function $\mu : [t_0, T] \to \mathbb{R}^D$ and matrix-valued covariance function $C : [t_0, T]^2 \to \mathbb{R}^{D \times D}$. Here the mean components are

$$\begin{aligned}\mu_j(t) &= e^{k_j(t_0 - t)} \langle x_j(t_0) \rangle \\ &+ \sum_{i \in \mathcal{I}} \frac{k_j}{k_j - a_j} \zeta_{i,j}^{\mathfrak{a}} (e^{(k_j - a_j)(\min\{t_i, t\} - t)} - e^{(k_j - a_j)(\max\{t_{i-1}, t_0\} - t)})\end{aligned}$$

and the covariance matrix function is $C(s,t) = diag(cov_{11}(s,t), ..., cov_{DD}(s,t))$ where

$$\begin{aligned}cov_{jj}(s,t) &= e^{(a_j - k_j)(t + s - 2t_0)} (\langle x_j^2(t_0) \rangle - \langle x_j(t_0) \rangle^2) \\ &+ \frac{\nu_j}{2(k_j - a_j)} [e^{(a_j - k_j)|t - s|} - e^{(k_j - a_j)(2t_0 - (s+t))}].\end{aligned}$$

B Deriving Lipschitz numbers

So far, we have established how knowledge of the Lipschitz number of a Lipschitz function can be utilized to exclude the presence of any negative function values on a compact domain based on a finite number of function evaluations. To employ this insight in the context of collision detection we will have to know a Lipschitz number L of the criterion functions $\gamma^{\mathfrak{a},\mathfrak{r}}$.

We may consider two cases:

- We have a belief quantified as a distribution over the smallest Lipschitz constant L'. Let the cumulative distribution function of this belief be denoted by $F : \mathbb{R} \to [0, 1]$. That is, $F(x) = \Pr[L' < x]$. If we desire a guaranteed success of collision detection of at least $\delta \in (0, 1)$ we invoke *Alg.*1 with a Lipschitz number $L \geq \min\{x \in \mathbb{R} | F(x) \geq \delta\}$ to detect non-positive values of $\gamma^{\mathfrak{a},\mathfrak{r}}$.

- If we do not have such a belief function or desire complete certainty in collision detection success it may be possible to derive a Lipschitz number for the $\gamma^{\mathfrak{a},\mathfrak{r}}$ based on the mean and covariance functions of the agent's stochastic trajectories. In turn, these may be derived from the agents' SDEs. How this can be accomplished is the subject of the remainder of this section.

B.1 Lipschitz arithmetic

In preparation of the derivations of Lipschitz numbers, we need to establish a few basic properties of Lipschitz continuous functions. As a convention, L_ϕ will always denote the Lipschitz constant of a Lipschitz continuous function ϕ.

Lemma B.1 (Lipschitz arithmetic)**.** *Let, $I, J \subset \mathbb{R}_+$. Let $f : \mathbb{R} \to \mathbb{R}$ be Lipschitz on I with Lipschitz number $L_I(f)$ and $g : \mathbb{R} \to \mathbb{R}$ be Lipschitz on J with Lipschitz number $L_J(g)$. We have:*

1. *Mapping $t \mapsto |f(t)|$ is Lipschitz on I with constant $L_I(f)$.*
2. *If g is Lipschitz on all of $J = f(I)$ the concatenation $g \circ f : t \mapsto g(f(t))$ is Lipschitz on I with Lipschitz constant $L_I(g \circ f) \leq L_J(g)\, L_I(f)$.*
3. *Let $r \in \mathbb{R}$. $r\, f : x \mapsto r\, f(x)$ is Lipschitz on I having a Lipschitz constant $L_I(r\, f) = |r|\, L_I(f)$.*
4. *$f + g : t \mapsto f(t) + g(t)$ has Lipschitz number at most $L_I(f) + L_J(g)$.*
5. *Let $m_f = \sup_{t\in I} f(t)$ and $m_g = \sup_{t\in I} g(t)$. Product function $f \cdot g : x \mapsto f(x)\, g(x)$ has a Lipschitz number on I which is at most $(m_f\, L_J(g) + m_g\, L_I(f))$.*
6. *Let $h(t) = \max\{f(t), g(t)\}, \forall t \in I \cap J$. We have $L_{I\cap J}(h) = \max\{L_I(f), L_J(g)\}$.*
7. *Let $b := \inf_{t\in I} |f(t)| > 0$ and let $\phi(t) = \frac{1}{f(t)}, \forall t \in I$. Then $L_I(\phi) \leq b^{-2}\, L_I(f)$ on I.*
8. *f cont. differentiable on $I \Rightarrow L_I(f) = \sup_{t\in I} |\dot{f}(t)|$.*
9. *Let $c \in \mathbb{R}$, $f(t) = c, \forall t \in I$. Then $L_I(f) = 0$.*
10. *$L_I(f^2) \leq 2\, s(f)\, L_I(f)$.*
11. *f cont. differentiable $\Rightarrow \forall q \in \mathbb{Q} : L_I(f^q) = |q| \sup_{\tau\in I} |f^{q-1}(\tau)\, \dot{f}(\tau)|$.*

Proof. **1)** We show $|f|$ has the same Lipschitz number as f. Let $t, t' \in I$ arbitrary. We enter a case differentiation:

1st case: $f(t), f(t') \geq 0$.

Hence, $\big||f(t)| - |f(t')|\big| = \big|f(t) - f(t')\big| \overset{f\ Lipschitz}{\leq} L_I(f)|t - t'|$.

2nd case: $f(t) \geq 0, f(t') \leq 0$.

Note, $|y| = -y$, iff $y \leq 0$. Hence, $\big||f(t)| - |f(t')|\big| \leq \big||f(t)| + |f(t')|\big| = \big||f(t)| - f(t')\big| = \big|f(t) - f(t')\big| \leq L_I(f)\,|t - t'|$.

3rd case: $f(t) \leq 0, f(t') \geq 0$. Completely analogous to 2nd case.

4th case: $f(t), f(t') \leq 0$.

$\big||f(t)| - |f(t')|\big| = \big|f(t') - f(t)\big| = \big|f(t) - f(t')\big| \overset{f\ Lipschitz}{\leq} L_I(f)|t - t'|$.

2) For arbitrary $t, t' \in I$ we have:

$\big|g(f(t)) - g(f(t'))|\big| \leq L_J(g)\,|f(t) - f(t')| \leq L_J(g)\,L_I(f)\,|t - t'|$ where the two inequalities are due to the Lipschitz properties of g and f, respectively.

3) For arbitrary $t, t' \in I, r \in \mathbb{R}$ we have:

$\big|r\,f(t) - r\,f(t')\big| = |r|\,|f(t) - f(t')| \leq |r|\,L_I(f)\,|t - t'|$.

4) For arbitrary $t, t' \in I, r \in \mathbb{R}$ we have:

$\big|g(t) + f(t) - (g(t') + f(t'))|\big| = \big|g(t) - g(t') + f(t) - f(t')|\big| \leq \big|g(t) - g(t')\big| + \big|f(t) - f(t')|\big| \leq (L_J(g) + L_I(f))\,|t - t'|$.

5) Let $t, t' \in I$, $d := f(t) - f(t')$.

$\big|g(t)f(t) - g(t')f(t')\big| = \big|g(t)(f(t') + d) - g(t')f(t')\big| = \big|\big(g(t) - g(t')\big)f(t') + g(t)d\big| \leq \big|g(t) - g(t')\big||f(t')| + \big|g(t)\big|\,|d| \leq L_I(g)|t - t'||f(t')| + \big|g(t)\big|\,L_I(f)|t - t'| \leq L_I(g)|t - t'|\,\sup_{t' \in I}\{|f(t')|\} + \sup_{t \in I}\{\big|g(t)\big|\}\,L_I(f)|t - t'| = \Big(L_I(g)\,\sup_{t' \in I}\{|f(t')|\} + \sup_{t \in I}\{\big|g(t)\big|\}\,L_I(f)\Big)|t - t'|$.

6) Proof in "' Nick Weaver, Lipschitz algebras"'.

7) Let $t, t' \in I$. $\big|\frac{1}{f(t)} - \frac{1}{f(t')}\big| = \big|\frac{f(t')}{f(t')f(t)} - \frac{f(t)}{f(t')f(t)}\big| = \frac{\big|f(t') - f(t)\big|}{\big|f(t')\big|\big|f(t)\big|} \leq \frac{L_I(f)|t - t'|}{\inf_{t \in I}|f(t)|}$.

8) Define $\ell := \sup_{t \in I}|\dot{f}(t)| = L_I(f)$. In two steps, we show that ℓ is the smallest Lipschitz constant.

Firstly, we show that it is a Lipschitz constant: Let $t, t' \in I, t < t'$. Due to the mean value theorem $\exists \xi \in [t, t'] \subset I : \frac{|f(t) - f(t')|}{|t - t'|} = |\dot{f}(\xi)| \leq \ell$. Secondly, we show that ℓ is the smallest Lipschitz constant: Let $\bar{\ell}$ be another Lipschitz constant such that $\bar{\ell} \leq \ell$. Since I is compact and $\dot{f}$ is continuous there is some $\xi \in I$ such that $\dot{f}(\xi) = \ell$. Pick any sequence $(t_k)_{k=1}^{\infty}$ such that $t_k \overset{k \to \infty}{\longrightarrow} \xi$. $\forall k : t_k \in I$ and $\bar{\ell}$ is a Lipschitz constant on I. Hence, $\bar{\ell} \geq \frac{|f(t_k) - f(\xi)|}{|t_k - \xi|} \overset{k \to \infty}{\longrightarrow} |\dot{f}(\xi)| = \ell$. Thus, $\bar{\ell} = \ell$.

9) Immediate consequence of 8).

10) $L(f^2) = L(f\,f) \leq s(f)L(f) + L(f)s(f)$ where the last inequality follows from property 5).

11) $L(f^q) \overset{8)}{=} \sup_{\tau\in I} |\frac{\mathrm{d}}{\mathrm{d}\,t} f^q(\tau)| = |q|\, \sup_{\tau\in I} |f^{q-1}(\tau)\, \dot{f}(\tau)|$.

□

Notice, that several of the inequalities in the Lemma are not tight (Eg. Ineq. (5), (10)). Therefore, it may sometimes be better not to apply it if instead one is able to determine the Lipschitz constant directly to yield a lower Lipschitz number.

Lemma B.2. *For $0 < a < b$ let $J \subset \mathbb{R}_+$ be the domain interval of square root function $\sqrt{\cdot} : J \to \mathbb{R}_+, t \mapsto \sqrt{t}$ such that $\inf J = a$ and $\sup J = b$. We have $L_J(\sqrt{\cdot}) \leq \frac{1}{2\sqrt{a}}$ where $L_J(\sqrt{\cdot})$ denotes the Lipschitz constant of the square root function on J.*

Proof. Applying Lem. B.1 and leveraging differentiability of the square root function reduces the problem of determining a Lipschitz constant to finding $\sup_{t\in J} |\frac{d\sqrt{\cdot}}{dt}(t)| = \sup_{t\in J} |\frac{1}{2\sqrt{t}}|$. Since $\frac{d^2}{dt^2}\sqrt{\cdot} = -\frac{1}{4\sqrt{\cdot}^3}$ attains only negative values on $J \subset \mathbb{R}_+$ we know that the first derivative $\frac{d\sqrt{\cdot}}{dt} = \frac{1}{2\sqrt{\cdot}}$ is strictly monotonously decreasing. Thus, $\sup_{t\in J} |\frac{d\sqrt{\cdot}}{dt}(t)| = \frac{d\sqrt{\cdot}}{dt}(\inf J) = \frac{1}{2\sqrt{\inf J}} = \frac{1}{2\sqrt{a}}$. □

B.2 An alternative collision-criterion function and the derivation of its Lipschitz number

Theorem B.3 (Alternative collision criterion function)**.** *Let spatial dimensionality $D = 2$. Let $\delta^{\mathfrak{a}} \in (0, 1)$ denote the maximum upper bound on instantaneous collision probability at time t agent $\mathfrak{a} \in \mathfrak{A}$ is allowed to tolerate. Let $\mu^{\mathfrak{a}}(t) \in \mathbb{R}^D$ be the mean of trajectory $x^{\mathfrak{a}}$ and $C^{\mathfrak{a}}_{ij}(t)$ be the spatial between dimensions i and j at time t. For $i \in \{1, 2\}, j \in \{1, 2\} - \{i\}$ let $r^{\mathfrak{a}}_i(t) := \sqrt{\frac{1}{2\delta^{\mathfrak{a}}}} \sqrt{C^{\mathfrak{a}}_{ii}(t) + \frac{\sqrt{C^{\mathfrak{a}}_{ii}(t)C^{\mathfrak{a}}_{jj}(t)(C^{\mathfrak{a}}_{ii}(t)C^{\mathfrak{a}}_{jj}(t) - (C^{\mathfrak{a}}_{ij}(t))^2)}}{C^{\mathfrak{a}}_{jj}(t)}}$.*

Let $\mathfrak{a}, \mathfrak{r}$ be two agents' plants whose radii sum to $\Lambda^{\mathfrak{a},\mathfrak{r}}$ with state trajectories $x^{\mathfrak{a}}, x^{\mathfrak{r}}$, respectively. Define

$$b^{\mathfrak{a},\mathfrak{r}}_j(t, t') := r^{\mathfrak{a}}_j(t) + r^{\mathfrak{r}}_j(t') + \Lambda^{\mathfrak{a},\mathfrak{r}}.$$

The function

$$\gamma^{\mathfrak{a},\mathfrak{r}}(t) := \max_{j\in\{1,\ldots,D\}} \{|\mu^{\mathfrak{a}}_j(t) - \mu^{\mathfrak{r}}_j(t)| - b^{\mathfrak{a},\mathfrak{r}}_j(t, t)\}$$

is a valid criterion function. That is, $\gamma^{\mathfrak{a},\mathfrak{r}}(t) > 0 \Rightarrow \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] < \delta^{\mathfrak{a}}$.

Proof. (Sketch) Let $t \in I$ be an arbitrary but fixed time. It is straight-forward to adapt the proof of Thm. 2 in [20] to our case showing that $\Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] < \delta^{\mathfrak{a}}$ if $|\mu^{\mathfrak{a}}_j(t) - \mu^{\mathfrak{r}}_j(t, t')| - b_j(t, t') > 0$ for at least one $j \in \{1, \ldots, D\}$. Hence, $\{t \in I | \gamma^{\mathfrak{a},\mathfrak{r}}(t, t) > 0\} \subset \{t \in I | \Pr[\mathfrak{C}^{\mathfrak{a},\mathfrak{r}}(t)] < \delta^{\mathfrak{a}}\}$. □

Definition B.4. *For future reference we define* $s : f \mapsto \sup_{t\in I} |f(t)|$ *and* $\iota : f \mapsto \inf_{t\in I} |f(t)|$ *on the space of continuous functions on interval* I*.*

Theorem B.5. *Let* $D = 2$ *be the dimensionality of state space. For any agent* $\mathfrak{a}$ *let* $\mu_j^{\mathfrak{a}} : I = [t_0, t_f] \to \mathbb{R}$ *denote the* j*th component of* $\mathfrak{a}$*'s mean function and* $C_{ij}^{\mathfrak{a}}(t)$ *the covariance of agent* $\mathfrak{a}$*'s trajectory between dimension* j *and* i *at time* t*.*

For $\mathfrak{q} \in \{\mathfrak{a}, \mathfrak{r}\}$ *and all* $i, j \in \{1, \dots, D\}$ *assume the* $\mu_j^{\mathfrak{q}}, C_{ij}^{\mathfrak{q}}$ *are Lipschitz on* I *with Lipschitz numbers* $L(\mu_j^{\mathfrak{q}})$ *and Lipschitz numbers* $L(C_{ij}^{\mathfrak{q}})$*, respectively.*

Let $\gamma^{\mathfrak{a},\mathfrak{r}}$ *denote the collision criterion function between agents* $\mathfrak{a}$ *and* $\mathfrak{r}$ *as defined in Thm. B.3.*

Then $\gamma^{\mathfrak{a},\mathfrak{r}}$ *is Lipschitz on* I *with Lipschitz constant*

$$\begin{aligned} L(\gamma^{\mathfrak{a},\mathfrak{r}}) &\le \max_{j\in\{1,\dots,D\}} \{L(\mu_j^{\mathfrak{a}} - \mu_j^{\mathfrak{r}}) + L(b_j^{\mathfrak{a},\mathfrak{r}})\} \\ &\le \max_{j\in\{1,\dots,D\}} \{L(\mu_j^{\mathfrak{a}}) + L(\mu_j^{\mathfrak{r}}) + L(b_j^{\mathfrak{a},\mathfrak{r}})\} \end{aligned}$$

where $\gamma^{\mathfrak{a},\mathfrak{r}}, b^{\mathfrak{a},\mathfrak{r}}, \alpha^{\mathfrak{a},\mathfrak{r}}$ *are defined as in Def. B.3 and* $\forall i \in \{1, 2\}, j \in \{1, 2\} - \{i\}$ $\forall \mathfrak{q} \in \{\mathfrak{a}, \mathfrak{r}\}$ *we have:*

1. $L(b_j^{\mathfrak{a},\mathfrak{r}}) \le \frac{1}{2}\big(L(\alpha_j^{\mathfrak{a}}) + L(\alpha_j^{\mathfrak{r}})\big)$,
2. $L(\alpha_i^{\mathfrak{q}}) \le \sqrt{\frac{1}{2\delta^{\mathfrak{q}}}}\, \frac{1}{\iota(g_i)} L(g_i)$

   where $g_i(t) := C_{ii}^{\mathfrak{q}}(t) + \sqrt{(C_{ii}^{\mathfrak{q}})^2 - C_{ii}^{\mathfrak{q}} \frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}}$ *and*

   where (i) $L(g_i) \le L(C_{ii}^{\mathfrak{q}}) + Q\, L\Big((C_{ii}^{\mathfrak{q}})^2\Big) + Q\, s\big(\big(C_{ij}^{\mathfrak{q}}\big)^2\big) L\Big(\frac{C_{ij}^{\mathfrak{q}}}{C_{jj}^{\mathfrak{q}}}\Big) + Q\, L\big(\big(C_{ij}^{\mathfrak{q}}\big)^2\big) s\Big(\frac{C_{ij}^{\mathfrak{q}}}{C_{jj}^{\mathfrak{q}}}\Big)$, *and*

   (ii) $L(g_i) \le L(C_{ii}^{\mathfrak{q}}) + Q\, L\Big((C_{ii}^{\mathfrak{q}})^2\Big) + Q\, s(C_{ii}^{\mathfrak{q}}) L\Big(\frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}\Big) + Q\, L\big(C_{ii}^{\mathfrak{q}}\big) s\Big(\frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}\Big)$.

   Here, $Q = \inf_{t\in I}\Big((C_{ii}^{\mathfrak{q}}(t))^2 - C_{ii}^{\mathfrak{q}}(t) \frac{\left(C_{ij}^{\mathfrak{q}}(t)\right)^2}{C_{jj}^{\mathfrak{q}}(t)}\Big) = \inf_{t\in I}(g_i - C_{ii}^{\mathfrak{q}})^2(t)$.
3. *also,* $L(\alpha_j^{\mathfrak{q}}) \le \frac{1}{2\sqrt{a}} L\Big(\frac{C_{jj}^{\mathfrak{q}}}{C_{ii}^{\mathfrak{q}}}\Big) s(\alpha_i^{\mathfrak{q}}) + s\Big(\sqrt{\frac{C_{jj}^{\mathfrak{q}}}{C_{ii}^{\mathfrak{q}}}}\Big)\, L(\alpha_1^{\mathfrak{q}})$*, where* $a = \inf_{t\in I} \frac{C_{jj}^{\mathfrak{q}}(t)}{C_{jj}^{\mathfrak{q}}(t)}$,
4. $L\Big(\frac{C_{ij}^{\mathfrak{q}}}{C_{jj}^{\mathfrak{q}}}\Big) \le L(C_{jj}^{\mathfrak{q}}) \iota(C_{jj}^{\mathfrak{q}})^{-2} s(C_{ii}^{\mathfrak{q}}) + s(\frac{1}{C_{jj}^{\mathfrak{q}}})\, L(C_{ii}^{\mathfrak{q}})$ *if* $\iota(C_{jj}) = \inf_{t\in I} |C_{jj}(t)| > 0$,
5. *and similarly,* $L\Big(\frac{(C_{ij}^{\mathfrak{q}})^2}{C_{jj}^{\mathfrak{q}}}\Big) \le L(C_{jj}^{\mathfrak{q}}) \iota(C_{jj}^{\mathfrak{q}})^{-2} s((C_{ij}^{\mathfrak{q}})^2) + s(\frac{1}{C_{jj}^{\mathfrak{q}}})\, L((C_{ij}^{\mathfrak{q}})^2)$ *if* $\iota(C_{jj}^{\mathfrak{q}}) = \inf_{t\in I} |C_{jj}^{\mathfrak{q}}(t)| > 0$.
6. $\forall i, j \in \{1, 2\} : L\big((C_{ij}^{\mathfrak{q}})^2\big) \le 2\, s(C_{ij}^{\mathfrak{q}})\, L\big(C_{ij}^{\mathfrak{q}}\big)$.

Proof. The equalities follow from successively applying Lem. B.1 to the definitions of the parts of the criterion function given in Thm. B.3. In our derivations we will note which of the properties of the Lemma we utilized as a superscript above the inequality sign. We have

$L(\gamma^{\mathfrak{a},\mathfrak{r}})$
$\overset{def.}{=} L(\max_{j\in\{1,\dots,D\}}\{|\mu_j^{\mathfrak{a}}-\mu_j^{\mathfrak{r}}|-b_j^{\mathfrak{a},\mathfrak{r}}\})$
$\overset{(6)}{=} \max_{j\in\{1,\dots,D\}} L(|\mu_j^{\mathfrak{a}}-\mu_j^{\mathfrak{r}}|-b_j^{\mathfrak{a},\mathfrak{r}})$
$\leq \max_{j\in\{1,\dots,D\}} L(|\mu_j^{\mathfrak{a}}-\mu_j^{\mathfrak{r}}|)+L(b_j^{\mathfrak{a},\mathfrak{r}})$
$\overset{(1,3,4)}{\leq} \max_{j\in\{1,\dots,D\}} L(\mu_j^{\mathfrak{a}})+L(\mu_j^{\mathfrak{r}})+L(b_j^{\mathfrak{a},\mathfrak{r}}).$

Proof of Ineq. 1): By definition, $L(b_j^{\mathfrak{a},\mathfrak{r}}) = L(\frac{1}{2}(\alpha_j^{\mathfrak{a}}(t)+\alpha_j^{\mathfrak{r}}(t'))+\Delta^{\mathfrak{a},\mathfrak{r}})$
$\overset{(3,4)}{\leq} \frac{1}{2}(L((\alpha_j^{\mathfrak{a}})+L(\alpha_j^{\mathfrak{r}}))+L(\Delta^{\mathfrak{a},\mathfrak{r}}) \overset{(9)}{=} \frac{1}{2}(L((\alpha_j^{\mathfrak{a}})+L(\alpha_j^{\mathfrak{r}}))$.
Furthermore, to prove the remaining inequalities, assume $\mathfrak{q}\in\{\mathfrak{a},\mathfrak{r}\}$.

Proof of Ineq. 2): Let $g(t) := C_{ii}^{\mathfrak{q}}(t)+\sqrt{(C_{ii}^{\mathfrak{q}})^2 - C_{ii}^{\mathfrak{q}}\frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}}$. Then, $L(\alpha_1^{\mathfrak{q}}) \overset{(3)}{=} \sqrt{\frac{2}{\delta^{\mathfrak{a}}}}\,L(\sqrt{g}) \leq \sqrt{\frac{2}{\delta^{\mathfrak{q}}}}\,\frac{1}{2\iota(g)}L(g)$
where the last inequality follows from Lem. B.2 and Lem. B.1.(2) as before. Furthermore, $L(g) \overset{(4)}{\leq} L(C_{ii}^{\mathfrak{q}}) + L(\sqrt{(C_{ii}^{\mathfrak{q}})^2 - C_{ii}^{\mathfrak{q}}\frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}})$
$= L(C_{ii}^{\mathfrak{q}}) + L\Big(\sqrt{(C_{ii}^{\mathfrak{q}})^2 - C_{ii}^{\mathfrak{q}}\frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}}\Big)$
$= L(C_{ii}^{\mathfrak{q}}) + Q\,L\Big((C_{ii}^{\mathfrak{q}})^2 - C_{ii}^{\mathfrak{q}}\frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}\Big)$ where according to Lem. B.2, $Q = \iota\Big((C_{ii}^{\mathfrak{q}})^2 - C_{ii}^{\mathfrak{q}}\frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}\Big)$. Hence, $L(g) \leq L(C_{ii}^{\mathfrak{q}}) + Q\,L\Big((C_{ii}^{\mathfrak{q}})^2\Big) + Q\,L\Big(C_{ii}^{\mathfrak{q}}\frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}\Big)$.

Proof of Ineq. 3): We have $L(\alpha_j^{\mathfrak{q}}(t)) \overset{def}{=} L(\sqrt{\frac{C_{jj}^{\mathfrak{q}}}{C_{ii}^{\mathfrak{q}}}}\alpha_i^{\mathfrak{q}})$
$\overset{(5)}{\leq} s(\sqrt{\frac{C_{jj}^{\mathfrak{q}}}{C_{ii}^{\mathfrak{q}}}})L(\alpha_i^{\mathfrak{q}}) + L(\sqrt{\frac{C_{jj}^{\mathfrak{q}}}{C_{ii}^{\mathfrak{q}}}})s(\alpha_i^{\mathfrak{q}})$
$\overset{(2)}{\leq} s(\sqrt{\frac{C_{22}^{\mathfrak{q}}}{C_{ii}^{\mathfrak{q}}}})L(\alpha_i^{\mathfrak{q}}) + L_J(\sqrt{\cdot})L(\frac{C_{jj}^{\mathfrak{q}}}{C_{ii}^{\mathfrak{q}}})s(\alpha_i^{\mathfrak{q}})$ where $J = \frac{C_{jj}^{\mathfrak{q}}}{C_{ii}^{\mathfrak{q}}}(I)$. Inequality 2) now follows from applying Lem. B.2.

From here we can make two alternative derivations:
i) $L(g) \leq L(C_{11}^{\mathfrak{q}}) + Q\,L\Big((C_{11}^{\mathfrak{q}})^2\Big) + Q\,s\big((C_{12}^{\mathfrak{q}})^2\big)L\Big(\frac{C_{11}^{\mathfrak{q}}}{C_{22}^{\mathfrak{q}}}\Big) + Q\,L\big((C_{12}^{\mathfrak{q}})^2\big)s\Big(\frac{C_{11}^{\mathfrak{q}}}{C_{22}^{\mathfrak{q}}}\Big)$.
Alternatively, one can obtain:
ii) $L(g) \leq L(C_{11}^{\mathfrak{q}}) + Q\,L\Big((C_{11}^{\mathfrak{q}})^2\Big) + Q\,s(C_{11}^{\mathfrak{q}})L\Big(\frac{\left(C_{12}^{\mathfrak{q}}\right)^2}{C_{22}^{\mathfrak{q}}}\Big) + Q\,L\big(C_{11}^{\mathfrak{q}}\big)s\Big(\frac{\left(C_{12}^{\mathfrak{q}}\right)^2}{C_{22}^{\mathfrak{q}}}\Big)$.

Proof of Ineq. 4): $L\Big(\frac{C_{ii}^{\mathfrak{q}}}{C_{jj}^{\mathfrak{q}}}\Big) = L\Big(C_{ii}^{\mathfrak{q}}\,\frac{1}{C_{jj}^{\mathfrak{q}}}\Big) \overset{(5)}{\leq} L\Big(\frac{1}{C_{jj}^{\mathfrak{q}}}\Big)\,s\Big(C_{ii}^{\mathfrak{q}}\Big) + L\Big(C_{ii}^{\mathfrak{q}}\Big)\,s\Big(\frac{1}{C_{jj}^{\mathfrak{q}}}\Big)$

$\overset{(7)}{\leq}\ b^{-2}\, L\Big(C^{\mathfrak{q}}_{jj}\Big)\, s\Big(C^{\mathfrak{q}}_{ii}\Big) + L\Big(C^{\mathfrak{q}}_{ii}\Big)\, s\Big(\frac{1}{C^{\mathfrak{q}}_{jj}}\Big)$ where $b \in \mathbb{R}_+$ chosen such that $C^{\mathfrak{q}}_{jj}(I) \cap [-b, b] = \emptyset$ (we assume such a b exists). A valid choice certainly is $b := \inf_{t\in I} |C^{\mathfrak{q}}_{jj}(t)|$.

Proof of Ineq. 5): Completely, analogous to proof of 4).

Proof of Ineq. 6): Consequence of Lem. B.1.(10).

□

The theorem provides a recipe to find a Lipschitz bound for the collision criterion function given known Lipschitz numbers of the trajectories' means and spatial covariance mappings.

However, since most equalities are not tight one should attempt to determine Lipschitz numbers directly wherever possible rather than using the inequalities provided in Lem. B.1. For instance, if one can determine the best Lipschitz number for $L\big((C^{\mathfrak{q}}_{ij})^2\big)$ directly (e.g. by utilizing Lem B.1.11) this would normally yield a better Lipschitz constant than obtained by expanding into $2\, s(C^{\mathfrak{q}}_{ij})\, L\big(C^{\mathfrak{q}}_{ij}\big)$ due to application of Lem. B.1.6.

Examining the terms in the inequalities we notice the occurrence of suprema of covariances $s(C_{ij})$ or inverted covariances of the form $s(\frac{1}{C_{ii}})$. The latter requires non-vanishing uncertainty in our model. Furthermore, note, the need to evaluate know the extrema is not to burdensome as they can be rapidly found by pre-existing Lipschitz optimizers which are highly efficient. However, in many cases the optima are known a priori. For instance, if one knows that the uncertainty monotonously increases over time we have e.g. $s(C^{\mathfrak{q}}_{ij}) = C^{\mathfrak{q}}_{ij}(\inf I)$ and $s(\frac{1}{C^{\mathfrak{q}}_{ij}}) = C^{\mathfrak{q}}_{ij}(\sup I)$. Alternatively, the covariances may allow for an analytic closed-form solution of the extremum which may be analytically derived before run-time.

We will revisit these issues in Sec. B.3 where we examine a concrete application of the theorem to a multi-agent control scenario.

B.3 A Lipschitz number for the criterion function of our feedback-controlled agents

Let $\mathfrak{a} \in \mathfrak{A}$ be an agent with controlled plant dynamics given by the Ito-SDE

$$dx^{\mathfrak{a}} = K(\xi^{\mathfrak{a}} - x^{\mathfrak{a}})dt + B\, dW.$$

Here $x^{\mathfrak{a}}(t) \in \mathbb{R}^D$ denotes the agent $\mathfrak{a}'s$ state (e.g. location), $\xi^{\mathfrak{a}}(t) \in \mathbb{R}^D$ is the agent's setpoint signal at time $t \in I = [t_0, t_f]$. Furthermore, $K = \mathrm{diag}(k_1, ..., k_D) > 0$ is the controller's gain matrix and $B = \mathrm{diag}(\sqrt{\nu}_1, ..., \sqrt{\nu}_D)$ reflects the magnitude of the uncertainties (disturbances). Let uncertain start state $x^{\mathfrak{a}}(t_0)$ be a normally distributed random vector. Assume $\mathfrak{a}$'s plan is $p^{\mathfrak{a}} = \big((t^{\mathfrak{a}}_i, \zeta^{\mathfrak{a}}_i)\big)_{0=1}^{H^{\mathfrak{a}}}$ where each $\zeta^{\mathfrak{a}}_i = (\zeta^{\mathfrak{a}}_{i,j})_{j=1}^{D}$. Let $t \geq t_0$, $\underline{i} = \arg\max_i\{t_i \leq t_0\}$,

$\bar{i} = \arg\min_i\{t_i \geq t\}$ and $\mathcal{I} := \{i \in \{1, ..., H^{\mathfrak{a}}\} | \underline{i} < i \leq \bar{i}\}$. Furthermore, let $\xi_j^{\mathfrak{a}}$ denote the jth component of step-function reference signal $\xi^{\mathfrak{a}}$.

For ease of notation we will drop the agent superscripts throughout the remainder of this subsection. The solution to agent $\mathfrak{a}$'s SDE is a Gaussian process with vector-valued mean function $\mu^{\mathfrak{a}} : [t_0, t_f] \to \mathbb{R}^D$ and matrix-valued covariance function $C^{\mathfrak{a}} : [t_0, t_f]^2 \to \mathbb{R}^{D \times D}$. By applying Ito-calculus to the suitable expectations of the SDE we can show that we have

$$\mu_j^{\mathfrak{a}}(t) = e^{k_j^{\mathfrak{a}}(t_0 - t)} \langle x_j^{\mathfrak{a}}(t_0) \rangle + k_j^{\mathfrak{a}}\, e^{-k_j^{\mathfrak{a}} t} \int_{t_0}^{t} e^{k_j^{\mathfrak{a}} \tilde{t}} \xi_j^{\mathfrak{a}}(\tilde{t}) d\tilde{t}.$$

The covariance matrix function is $(s,t) \mapsto \mathrm{diag}(\mathrm{cov}_{11}^{\mathfrak{a}}(s,t), ..., \mathrm{cov}_{DD}^{\mathfrak{a}}(s,t))$ where

$$\begin{aligned}\mathrm{cov}_{jj}^{\mathfrak{a}}(s,t) = e^{-k_j^{\mathfrak{a}}(t+s-2t_0)}(\langle \big(x_j^{\mathfrak{a}}(t_0)\big)^2 \rangle - \langle x_j^{\mathfrak{a}}(t_0) \rangle^2) \\ + \frac{\nu_j^{\mathfrak{a}}}{2k_j^{\mathfrak{a}}}[e^{-k_j^{\mathfrak{a}}|t-s|} - e^{k_j^{\mathfrak{a}}(2t_0 - (s+t))}].\end{aligned}$$

Using the notation of Thm. B.5 we have
$C_{jj}^{\mathfrak{a}}(t) = \mathrm{cov}_{jj}^{\mathfrak{a}}(t,t)$

$$\begin{aligned}&= e^{-k_j^{\mathfrak{a}} 2(t-t_0)}(\langle \big(x_j^{\mathfrak{a}}(t_0)\big)^2 \rangle - \langle x_j^{\mathfrak{a}}(t_0) \rangle^2) + \frac{\nu_j^{\mathfrak{a}}}{2k_j^{\mathfrak{a}}}[1 - e^{k_j^{\mathfrak{a}} 2(t_0 - t)}] \\ &= e^{-k_j^{\mathfrak{a}} 2(t-t_0)}\big(C_{jj}^{\mathfrak{a}}(t_0) - \frac{\nu_j^{\mathfrak{a}}}{2k_j^{\mathfrak{a}}}\big) + \frac{\nu_j^{\mathfrak{a}}}{2k_j^{\mathfrak{a}}}.\end{aligned}$$

where $C_{jj}^{\mathfrak{a}}(t_0)$ is assumed to be a known quantification of the initial state uncertainty.

Next we will derive Lipschitz constants for the the component means and covariances which is necessary to derive a Lipschitz number for the collision criterion function.

Firstly, we consider the mean function. Defining $v^{\mathfrak{a}}(t) := e^{k_j^{\mathfrak{a}}(t_0 - t)} \langle x_j^{\mathfrak{a}}(t_0) \rangle$, $w^{\mathfrak{a}}(t) := k_j^{\mathfrak{a}}\, e^{-k_j^{\mathfrak{a}} t}$ and $\dot{q}^{\mathfrak{a}}(t) := e^{k_j^{\mathfrak{a}} t} \xi_j^{\mathfrak{a}}(t)$ we we can restate the component mean function $m_j^{\mathfrak{a}}$ as
$\mu_j^{\mathfrak{a}}(t) = v_j^{\mathfrak{a}}(t) + w_j^{\mathfrak{a}}(t)\, q_j^{\mathfrak{a}}(t)$.
Leveraging Lem. B.1 we see that

$$\begin{aligned}L\big(\mu_j^{\mathfrak{a}}\big) &\leq L(v_j^{\mathfrak{a}}) + s(w_j^{\mathfrak{a}})\, L(q_j^{\mathfrak{a}}) + s(q_j^{\mathfrak{a}})\, L(w_j^{\mathfrak{a}}) &(5)\\ &\leq s(\dot{v}_j^{\mathfrak{a}}) + s(w_j^{\mathfrak{a}})\, s(\dot{q}_j^{\mathfrak{a}}) + s(q_j^{\mathfrak{a}})\, s(w_j^{\mathfrak{a}}) &(6)\end{aligned}$$

where as before $s(f) = \sup_{t \in I} |f(t)|$ for any function f. Evaluation of the suprema depends on the setpoint signal ξ and on the k_j. For instance, choosing a constant setpoint ξ and $k_j > 0$ would yield:

- $s(\dot{v}_j^{\mathfrak{a}}) = \sup_{t \in [t_0, t_f]} | - k_j^{\mathfrak{a}} \langle x_j^{\mathfrak{a}}(t_0) \rangle e^{k_j^{\mathfrak{a}} t_0} e^{-k_j^{\mathfrak{a}} t}| = |k_j \langle x_j^{\mathfrak{a}}(t_0) \rangle|$ where the last equality holds since $e^{-k_j^{\mathfrak{a}} t}$ decreases monotonically,

- $s(w_j^{\mathfrak{a}}) = \sup_{t\in[t_0,t_f]} |k_j^{\mathfrak{a}} e^{-k_j^{\mathfrak{a}} t}| = |k_j^{\mathfrak{a}}| e^{-k_j^{\mathfrak{a}} t_0}$,
- $s(\dot{q}_j^{\mathfrak{a}}) = \sup_{t\in[t_0,t_f]} |\xi_j^{\mathfrak{a}}\, e^{k_j^{\mathfrak{a}} t}| = |\xi_j^{\mathfrak{a}}|\, e^{k_j^{\mathfrak{a}} t_f}$,
- $s(q_j^{\mathfrak{a}}) = \sup_{t\in[t_0,t_f]} |\int_{t_0}^{t} e^{k_j^{\mathfrak{a}}\tilde{t}} \xi_j^{\mathfrak{a}} d\tilde{t}| = |\frac{\xi_j^{\mathfrak{a}}}{k_j^{\mathfrak{a}}}[e^{k_j^{\mathfrak{a}} t_f} - e^{k_j^{\mathfrak{a}} t_0}]|$. Here we leveraged the monotonicity of the exponential function.

  Next, we derive Lipschitz constants for the covariances:

  Note the cross-covariances are zero $C_{ij}^{\mathfrak{a}}(t) = 0, \forall t, i \neq j$. Fortunately, the diagonal of the covariance matrix function are also continuously differentiable. In particular, we have $\dot{C}_{jj}^{\mathfrak{a}}(t) = -k_j^{\mathfrak{a}} 2 e^{-k_j^{\mathfrak{a}} 2(t-t_0)}\big(C_{jj}^{\mathfrak{a}}(t_0) - \frac{\nu_j^{\mathfrak{a}}}{2k_j^{\mathfrak{a}}}\big)$. We can once again utilize Lem. B.1 yielding a Lipschitz bound

$$\begin{aligned} L(C_{jj}^{\mathfrak{a}}) &\leq s\big(\dot{C}_{jj}^{\mathfrak{a}}\big) \\ &\leq \sup_{t\in[t_0,t_f]} |k_j^{\mathfrak{a}} 2\big(C_{jj}^{\mathfrak{a}}(t_0) - \frac{\nu_j^{\mathfrak{a}}}{2k_j^{\mathfrak{a}}}\big)| e^{-k_j^{\mathfrak{a}} 2(t-t_0)} \\ &= |k_j^{\mathfrak{a}} 2\big(C_{jj}^{\mathfrak{a}}(t_0) - \frac{\nu_j^{\mathfrak{a}}}{2k_j^{\mathfrak{a}}}\big)|. \end{aligned}$$

  where the last equality follows from the fact that $t \mapsto \exp(-k_j 2(t-t_0))$ is monotonically decreasing.

In summary, we have found

$$L(C_{jj}^{\mathfrak{a}}) \leq |k_j^{\mathfrak{a}} 2\big(C_{jj}^{\mathfrak{a}}(t_0) - \frac{\nu_j^{\mathfrak{a}}}{2k_j^{\mathfrak{a}}}\big)|, \tag{7}$$

$$L(C_{12}^{\mathfrak{a}}) = L(C_{21}^{\mathfrak{a}}) = 0, \tag{8}$$

$$L(\mu_j^{\mathfrak{a}}) \leq s(\dot{v}_j^{\mathfrak{a}}) + s(w_j^{\mathfrak{a}}) s(\dot{q}_j^{\mathfrak{a}}) + s(q_j^{\mathfrak{a}}) s(w_j^{\mathfrak{a}}), \tag{9}$$

$$= |k_j^{\mathfrak{a}} \langle x_j^{\mathfrak{a}}(t_0)\rangle| + |k_j^{\mathfrak{a}}||\xi_j^{\mathfrak{a}}|\, e^{k_j^{\mathfrak{a}}(t_f - t_0)} + |\xi_j^{\mathfrak{a}}\, [e^{k_j^{\mathfrak{a}}(t_f-t_0)} - 1]|. \tag{10}$$

Next, we combine our estimates of the mean and covariances with Lem. B.1 and Thm. B.5 to derive a Lipschitz number for the criterion function defined in Thm. B.3. Let $\mathfrak{q} \in \{\mathfrak{a}, \mathfrak{r}\}$. Since $L(C_{12}^{\mathfrak{q}}) = 0$ we have $g_i(t) := C_{ii}^{\mathfrak{q}}(t) + \sqrt{(C_{ii}^{\mathfrak{q}})^2 - C_{ii}^{\mathfrak{q}} \frac{\left(C_{ij}^{\mathfrak{q}}\right)^2}{C_{jj}^{\mathfrak{q}}}} = 2C_{ii}^{\mathfrak{q}}(t)$. By Thm. B.5.2, this implies $L(\alpha_i^{\mathfrak{q}}) \leq \sqrt{\frac{1}{2\delta^{\mathfrak{q}}}} \frac{1}{\iota(g_i)} L(g_i) = \sqrt{\frac{1}{2\delta^{\mathfrak{q}}}} \frac{1}{\iota(2C_{ii}^{\mathfrak{q}})} L(2C_{ii}^{\mathfrak{q}}) = \sqrt{\frac{1}{2\delta^{\mathfrak{q}}}} \frac{1}{\iota(C_{ii}^{\mathfrak{q}})} L(C_{ii}^{\mathfrak{q}})$ where the last equality is due to Lem. B.1.3 and due to the fact that $\inf_t |r\, f(t)| = |r| \inf_t |f(t)|$ for all constants r, functions f. Next, we determine $\iota(C_{ii}^{\mathfrak{q}})$. By inspecting its derivative, we notice that $C_{ii}^{\mathfrak{q}}$ is strictly monotonously increasing iff $C_{ii}^{\mathfrak{q}}(t_0) - \frac{v_i^{\mathfrak{q}}}{2k_i^{\mathfrak{q}}} < 0$ and monotonously decreasing otherwise. Also, $C_{ii}^{\mathfrak{q}}$ does not attain negative values implying $C_{ii}^{\mathfrak{q}} = \big|C_{ii}^{\mathfrak{q}}\big|$.

Hence, $\iota(C_{ii}^{\mathfrak{q}}) = \inf\{|C_{ii}^{\mathfrak{q}}(I)|\} = \inf\{C_{ii}^{\mathfrak{q}}(I)\}$

$= \begin{cases} C_{ii}^{\mathfrak{q}}(t_0), \text{ if } C_{ii}^{\mathfrak{q}}(t_0) < \frac{v_i^{\mathfrak{q}}}{2k^{\mathfrak{q}}}; \\ C_{ii}^{\mathfrak{q}}(t_f), \text{ if } C_{ii}^{\mathfrak{q}}(t_0) \geq \frac{v_i^{\mathfrak{q}}}{2k_i^{\mathfrak{q}}}; \end{cases}$ Now, we have all the necessary ingredients to utilize Thm. B.5 in order to wrap-up:

$L(\gamma^{\mathfrak{a},\mathfrak{r}}) \leq \max_j\{L(\mu_j^{\mathfrak{a}}) + L(\mu_j^{\mathfrak{r}}) + L(b_j^{\mathfrak{a},\mathfrak{r}})\}$
$\leq \max_j\{L(\mu_j^{\mathfrak{a}}) + L(\mu_j^{\mathfrak{r}}) + \frac{1}{2}L(\alpha_j^{\mathfrak{a}}) + \frac{1}{2}L(\alpha_j^{\mathfrak{r}})\}$
$\leq \max_j\big\{|k_j^{\mathfrak{a}}\mu_j^{\mathfrak{a}}(t_0)| + |k_j^{\mathfrak{a}}||\xi_j^{\mathfrak{a}}|e^{k_j^{\mathfrak{a}}(t_f-t_0)} + |\xi_j^{\mathfrak{a}}\left[e^{k_j^{\mathfrak{a}}(t_f-t_0)} - 1\right]| + |k_j^{\mathfrak{r}}\mu_j^{\mathfrak{r}}(t_0)| +$
$|k_j^{\mathfrak{r}}||\xi_j^{\mathfrak{r}}|e^{k_j^{\mathfrak{r}}(t_f-t_0)} + |\xi_j^{\mathfrak{r}}\left[e^{k_j^{\mathfrak{r}}(t_f-t_0)} - 1\right]| + \frac{1}{2}\sqrt{\frac{1}{2\delta^{\mathfrak{a}}}}\frac{1}{\iota(C_{jj}^{\mathfrak{a}})}L(C_{jj}^{\mathfrak{a}}) + \frac{1}{2}\sqrt{\frac{1}{2\delta^{\mathfrak{r}}}}\frac{1}{\iota(C_{jj}^{\mathfrak{r}})}L(C_{jj}^{\mathfrak{r}})\big\}$.

We can see that this Lipschitz number might adopt large values in certain parts of the domain. Therefore, it might be helpful to recompute the Lipschitz numbers adaptively for different parts of the domain.

C Utilising priors encoding belief over change of sign of a criterion function

Detection of collisions is based on excluding the possibility of negative criterion function values. However, as these functions are non-convex any numerical procedure executed on a digital computer has to achieve this with only a finite number of function evaluations. Given this, what is our confidence in not having missed a negative criterion function value?

Thus far, we have proposed using a knowledge (i.e. a prior) about a Lipschitz number of the criterion function to rule out collisions in continuous time based on a finite number of samples. In addition to the Lipschitz-based method presented above, we will now consider an alternative method that assumes a prior belief about the anticipated change of sign of a criterion function.

Before commencing it will prove helpful to introduce the notion of a *sign change point (SCP)*. An SCP is a time step which is the border between two changes in sign of a function. More precisely, time t is an SCP of function f if there exist open intervals $I' := (t', t)$ and $I'' := (t, t'')$ such that $\text{sgn}(f(\tau')) \neq \text{sgn}(f(\tau'')), \forall \tau' \in I', \tau'' \in I''$.

To give an example, consider the function

$$\phi : t \mapsto -t\,\chi_{\mathbb{R}_-}(t) + 0\,\chi_{[0,1]}(t) + (t-1)\chi_{\mathbb{R}_{>1}}(t).$$

As before, χ_S denotes the indicator function of set S. Function ϕ has exactly two SCPs – at $t = 0$ and $t = 1$.

Resuming with our discussion, assume we are given $f(t_0), ..., f(t_k)$ on a lattice of times ($0 = t_0 < ... < t_k = T$). If $f(t_i) \leq 0$ for some t_i we will want to conservatively assume a collision has occurred. On the other hand, if all evaluations are positive we desire to specify our confidence that all intermittent unobserved values are. This is the case if no SCPs occur. The presence of an odd number of SCPs between two time steps t_i, t_{i+1} is detectable by checking $\text{sgn}(f(t_i)) \neq \text{sgn}(f(t_{i+1}))$. In fact, if the total number n of SCPs in $[0, T]$ is

odd we will detect a change of sign. By contrast, if an even number of SCPs occur we have $\text{sgn}(f(t_i)) = \text{sgn}(f(t_{i+1}))$ and hence, will be oblivious of negative function values in the interval (t_i, t_{i+1}).

Now, assume we are given a distribution $Q : \mathbb{N} \to [0, 1]$ representing our belief over the number of occurring SCPs. By the law of total probability, our belief that we will miss the existence of a collision is

$$\sum_{n \in 2\mathbb{N}} P_{n,k} Q[n] \tag{11}$$

where $P_{n,k}$ denotes the probability of missing the existence of a collision during collision detection given that n SCPs occur in the interior of the lattice.

In preparation of the next theorem we need the following result:

Lemma C.1 (Improved bound)**.** *Given a set $S = \{s_1, ..., s_n\}$ of $n \in 2\mathbb{N}$ objects let $\mathcal{P}$ be the number of ways the set can be partitioned into pairs, i.e. $\mathcal{P} = |\{\{\mathfrak{P}_1, ..., \mathfrak{P}_{n/2}\} | \forall i \neq j : \mathfrak{P}_i \cap \mathfrak{P}_j = \emptyset, \bigcup_i \mathfrak{P}_i = \{s_1, ..., s_n\}, \forall i \exists q \neq r : \mathfrak{P}_i = \{s_q, s_r\}\}|$.*

We have $\mathcal{P} \leq \binom{n(n-1)/2}{n/2}$.

Proof. We can create $(n^2 - n)/2 = n(n-1)/2$ distinct sets of the form $\{s_q, s_r\}$ of cardinality two. That is, $|T| = n(n-1)/2$ where $T = \{\{s_i, s_j\} | i \neq j, i, j = 1, ..., n\}$.

To generate a partition $\{\mathfrak{P}_1, ..., \mathfrak{P}_{n/2}\}$ we need to select a subset of T containing $n/2$ two-element sets. Conservatively (not taking into account that not every n/2 -element subset is an actual partition), this could be done in at most $\binom{n(n-1)/2}{n/2}$ ways. Hence $\mathcal{P} \leq \binom{n(n-1)/2}{n/2}$.

□

Theorem C.2. *Assume we are given $n \in 2\mathbb{N}_0$ SCPs $s_1, ..., s_n$ whose locations are drawn independently from an identical distribution (drawn i.i.d.). Furthermore, we are given a grid of test points $0 = t_0 < t_1 < ... < t_k = T$ where the intermediate times are chosen such that $\forall i \in \{1, ..., n\}, j \in \{1, ..., k\} : \Pr[s_i \in (t_{j-1}, t_j)] = 1/k$. The probability of missing the existence of a collision by looking for non-positive elements in the sample $f(t_0), ..., f(t_k)$ is:*

$$P_{n,k} \leq \frac{\mathcal{P}}{\sqrt{k^n}} \leq \frac{\binom{n(n-1)/2}{n/2}}{\sqrt{k^n}}$$

where $\mathcal{P}$ is a function of n (but not of k) as defined in Lemma C.1. In particular, we have $\lim_{k \to \infty} P_{n,k} = 0$.

Proof. We define the sample space $\Omega := \{(b_1, ..., b_n) \in \{1, ..., k\}^n\}$ where each $b_i \in \{1, ..., k\}$ denotes the index of the time interval ("bin") $(t_{b_i - 1}, t_{b_i}]$ the ith SCP s_i falls into $(i = 1, ..., n)$. Due to the assumption that the assignment of SCP to bin is i.i.d., each sample has equal probability and we can compute $P_{n,k}$ as a Laplace probability. That is, $P_{n,k} = \frac{|G|}{|\Omega|}$ where G is the set of events

describing that no bin contains an odd number of SCPs (because if at least one does contain an odd number we detect the presence of a collision).

Obviously, $|\Omega| = k^n$.

On the other hand, $G = \{(v_1, ..., v_k) | \sum_{j=1}^{k} v_j = n, \forall j \in \{1, ..., k\} : v_j \in \{0, ..., n\} \cap 2\mathbb{N}_0\}$ where $v_j \in \{0, ..., n\}$ denotes the number of SCPs falling into bin $j \in \{1, ..., k\}$. We will find an upper bound on $G's$ cardinality by constructing a finite set H for whose cardinality one can easily establish $\sqrt{k}^n \mathcal{P}$ as an upper bound. We show that one can define an injective function $\psi : G \to H$. The latter establishes that $|G| \leq |H|$. Thus, $\frac{|G|}{|\Omega|} \leq \frac{|H|}{|\Omega|} \leq \frac{\mathcal{P}}{\sqrt{k}^n}$ which will hence conclude the proof.

We generate H by invoking a two-stage process (where in each stage it is easy to enumerate all possible elements that are generated). In the first stage, we partition the SCPs into $n/2$ pairs (which we always can since we assumed n to be even). In the second stage, we assign these pairs to the bins in which the pairs are merged into joint sets of SCPs. Therefore, $H = \{(M_1, ..., M_k) | M_1 \subset \{s_1, ..., s_n\}, ..., M_k \subset \{s_1, ..., s_n\}, |M_1|, ..., |M_k| \in 2\mathbb{N}_0, |M_1| + ... + |M_k| = n, H = \bigcup_i M_i\}$.

Let $\mathcal{P}$ be the number of ways in which one can partition the n SCPs into $n/2$ (unordered) pairs i.e. $\mathcal{P} = |\{\{\mathfrak{P}_1, ..., \mathfrak{P}_{n/2}\} | \forall i \neq j : \mathfrak{P}_i \cap \mathfrak{P}_j = \emptyset, \bigcup_i \mathfrak{P}_i = \{s_1, ..., s_n\}, \forall i \exists q \neq r : \mathfrak{P}_i = \{s_q, s_r\}\}|$. (Cf. Lem. C.1 for a bound). In the second stage, the pairs are distributed among the k bins (intervals) (which can be done in $\mathcal{B} = k^{n/2}$ ways) before the sets within each bin are merged. The number of different paths the process can take to generate an element in H is $\mathcal{P}$ (number of partitions into pairs) multiplied with $\mathcal{B}$ (number of ways the pairs constituting the partition can then be distributed into the bins).

By construction, each final assignment (subsets of SCPs to bins) generated by the two-stage process is an element of H. Conversely, let $(M_1, ..., M_k) \in H$ then it is easy to verify it could be generated by the two-stage process (however, there may be multiple paths in the process generating the same element of H). Hence, $|H| \leq \mathcal{P}\mathcal{B}$.

We finalize our considerations by defining the function $\psi : G \to H, (v_1, ..., v_n) \mapsto (M_1, ..., M_k)$ where $M_1 := \{s_1, ..., s_{v_1}\}$ and for $i > 1$: $M_i := \{s_{1+w_i}, ..., s_{v_i+w_i}\}$ where $w_i = \sum_{j<i} v_j$. It is easy to see that ψ maps different $(v_1, ..., v_n) \in G$ to different $(M_1, ..., M_k) \in H$. Hence, ψ is injective. Since both H and G are finite this implies $|G| \leq |H|$. Wrapping up, $\frac{|G|}{|\Omega|} \leq \frac{|H|}{|\Omega|} \leq \frac{\mathcal{P}\mathcal{B}}{k^n} \leq \frac{\mathcal{P}k^{n/2}}{k^n} = \frac{\mathcal{P}}{k^{n/2}} \overset{Lem.C.1}{\leq} \frac{\binom{n(n-1)/2}{n/2}}{\sqrt{k}^n}$. □

In conjunction with Eq. 11, Thm. C.2 can provides a recipe how to do collision detection based on a finite number of criterion function samples such that our confidence in not having overlooked an existing collision is above a certain threshold θ. To this end we require a prior Q over the number of SCPs. Let N be a random variable quantifying the number of occurring SCPs and let M denote the event that we would not detect the existence of an SCP.

By the law of total probability we have

$$\Pr[M|k] = \sum_{n\in\mathbb{N}_0} \Pr[M|N=n,k]\, Q[N=n] \tag{12}$$

$$= \sum_{n\in 2\mathbb{N}_0} \Pr[M|N=n,k]\, Q[N=n] \tag{13}$$

$$= \sum_{n\in 2\mathbb{N}_0} P_{n,k}\, Q[N=n] \tag{14}$$

where the last equality follows from $\Pr[M|N=n,k] = 0, (n \in 2\mathbb{N}-1), k \geq 1$.

Given a threshold $\theta \in (0,1)$ we can then utilize Thm. C.2 to choose a lattice resolution k such that

$$\Pr[M|k] \overset{!}{<} \theta.$$

To illustrate this we provide to simple examples:

Example C.3 (Finite-support prior). As a simplistic example, assume we desire to detect a collision between two agents $\mathfrak{a}, \mathfrak{r}$ with plans containing two setpoints each (stabilizing their start and goal state each) and with linear dynamics. We know that at their start and goal locations, no collisions occur. Given this our belief over the number n of SCPs may be $Q[N=0] = 0.5$ $Q[N=1] = .1, Q[N=2] = .4$. Furthermore, in the given time interval [0,T] we assume a flat prior distribution giving rise to an equidistant lattice $0 = t_0 < t_1 < ... < t_k = T$ of samples during collision detection. That is, $t_{i+1} - t_i = T/k$ for some k. Our collision detection method now inspects our criterion function values $\Gamma^{\mathfrak{a},\mathfrak{r}}(t_0), ..., \Gamma^{\mathfrak{a},\mathfrak{r}}(t_k)$. If all values are positive we assume no collision has occurred. How large should we set k in order to ensure that our confidence in this assertion is at least $\theta \in (0,1)$? By Thm. C.2, we know that the probability of having missed a collision with this simple method is less than $\frac{0.4}{\sqrt{k^2}}$. Therefore, we need to set k such that $1 - \frac{0.4}{\sqrt{k^2}} \geq \theta$ which is equivalent to setting the number k of criterion function evaluations to $k \geq \frac{0.4}{1-\theta}$.